Shared Intelligence and Misaligned Firm Entry

Feng Chi[*] Nathan Yang[†]

July 2, 2026

Abstract

Firms increasingly rely on shared intelligence platforms for market entry. While these systems improve forecast accuracy, they also route decisions through a common processing layer that creates correlated errors unrelated to fundamentals. We develop a Bayesian entry game in which firms choose between proprietary research and a shared platform. Platform adoption generates two externalities: synchronized entry and correlated screening errors. A planner who internalizes these costs always prefers less adoption than one who values only firm profits. This external wedge is independent of whether entry decisions are strategic complements or substitutes, and over-adoption arises across a broad range of environments.

Keywords: Firm Entry; Correlated Errors; Information Externalities; Social Value of Information; Algorithmic Monoculture; Analytics Vendor Concentration; Economics of Data

[*]University of Illinois Urbana-Champaign, Department of Economics. Email: fchi@illinois.edu.
[†]University of Illinois Urbana-Champaign, Gies College of Business. Email: ncyang@illinois.edu.

1 Introduction

Retail entry is increasingly shaped by a common set of algorithmic forecasting platforms. Historically, firms relied on local, firm-specific information (Love et al. 1986). A regional grocer learned its trade area through direct experience, while a restaurant chain relied on proprietary market research. Recent work shows that the quality and timing of information can materially shape firm entry and survival (Chi 2025). We study a related but distinct margin: not whether firms have more timely information, but whether many firms process information through the same forecasting layer. Today, national chains, regional grocers, and franchisees increasingly rely on the same external analytics providers (e.g., CoStar, Placer.ai, and Buxton) to forecast trade-area demand, foot traffic, and competitive density.[1] These platforms can improve forecast accuracy by pooling data and reducing firm-specific noise. They also route many firms' decisions through a common processing layer. If that layer contains errors unrelated to local fundamentals, firms may make mistakes that are not only inaccurate but correlated.

This possibility is readily visible in the quick-service restaurant industry. In county-level data, McDonald's and Burger King exhibit strong geographic alignment: counties that appear attractive to one chain also tend to appear attractive to the other, even after accounting for observable demand fundamentals such as population, density, and income. This alignment is especially pronounced in counties with greater exposure to analysis using AI (Felten et al. 2021). Such co-location may reflect common demand fundamentals, agglomeration economies, or strategic imitation. The rapid proliferation of shared analytics vendors suggests an additional possibility. Firms may enter the same markets not only because those markets are attractive, but also because the same forecasting system points them there.

Standard models of retail entry do not readily generate this form of correlation.[2] In the existing

[1] Industry evidence suggests that a small number of providers are widely used across retail and commercial real estate. For example, Esri reports that seven of the ten largest U.S. retailers rely on its geospatial analytics platform, while firms such as Placer.ai provide location intelligence services to clients spanning retail, real estate, and hospitality sectors. These platforms are also frequently deployed through intermediaries such as commercial real estate advisory firms, implying that many decision-makers rely on overlapping data sources and processing systems. See `https://www.esri.com/en-us/industries/retail/overview` and `https://www.placer.ai/company/customers`.

[2] In canonical frameworks (Bresnahan and Reiss 1991, Seim 2006, Jia 2008, Holmes 2011), firms receive information about profitability, and correlation arises either from strategic interaction or from information about a common underlying state. In these cases, correlation reflects economically relevant information, and forecast errors are largely independent across firms. Forecasts from platforms differ because adoption does not make the common component more informative; instead, the same error is propa-

literature, correlation in firm behavior is typically attributed to strategic interaction, reputational herding (Scharfstein and Stein 1990, Froot et al. 1992), or informational cascades (Banerjee 1992, Bikhchandani et al. 1992). These mechanisms generate correlation because firms learn from common information or from one another's actions. By contrast, the channel we study involves no learning from others. Firms act on a common error embedded in a shared forecasting system, so correlation arises even when each firm conditions only on its own forecast and no firm observes another's action. Recent evidence suggests that algorithmic prediction can improve individual decisions while increasing correlation across firms. Firms investing heavily in AI exhibit greater systematic risk (Babina et al. 2024), and firms sharing data through common software development kits exhibit co-movement that survives controls for product similarity (Bian et al. 2024). What remains less understood is how welfare changes when many firms rely on the same forecasting infrastructure.

This distinction also separates our mechanism from related work on the productivity gains from shared data and forecasting infrastructure. As Farboodi and Veldkamp (2023) emphasize, data affect markets by improving prediction, sharpening targeting, changing competition, and altering the value of firms and financial assets. In models of data-driven automation, common data raise factor productivity and propagate across tasks, making shared infrastructure an engine of growth rather than a source of correlated error (Farboodi et al. 2026, Farboodi and Veldkamp 2021, Jones and Tonetti 2020). In those settings, whether shared data are accumulated too quickly or too slowly depends on how substitutable the affected activities are. Our focus is different. We isolate the part of a shared signal that remains after conditioning on fundamentals and therefore carries no productive content. This non-informative component generates correlated errors across adopters, and the resulting externality always points toward less platform adoption. The welfare wedge we study arises from correlated errors rather than from the productive component of shared data. It has a precedent in data markets: Acemoglu et al. (2022) show that when one user's data is informative about others, individually rational data sharing depresses the price of data and imposes an externality that markets do not internalize; our shared processing layer transposes that logic from the supply of data to the consumption of a common inference technology. Veldkamp and Chung (2024) survey how the nonrivalry and increasing returns of data reshape aggregate outcomes, a mechanism that here operates through the correlation, rather than the quantity, of firms' information.

gated across more decisions.

We develop a model in which firms choose between proprietary research and a shared forecasting platform. Proprietary research generates firm-specific forecast errors, whereas the shared platform reduces firm-specific noise but introduces a common processing error across adopters. This tradeoff creates a gap between privately and socially optimal adoption. Platform adoption generates two external costs that firms do not internalize: synchronized entry that affects parties exposed to aggregate fluctuations in commercial activity, and correlated screening errors that affect locations evaluated through the same forecasting technology. A planner who internalizes either cost chooses strictly less adoption than a planner who values only firm profits.

A central result is that this external wedge does not depend on whether entry decisions are strategic complements or substitutes. For the location-level welfare cost, the size of the wedge depends only on the adoption rate and the signal structure. The shared processing externality therefore differs from classical entry distortions, whose direction depends on the nature of strategic interaction. The shared processing externality always points toward less adoption.

The shared processing externality operates alongside the classical entry externality of Mankiw and Whinston (1986) and Spence (1976). Because firms ignore their effect on aggregate entry, decentralized outcomes are already inefficient even in the absence of shared processing. Whether the market ultimately over- or under-adopts therefore depends on the interaction between the classical entry distortion and the new externality introduced by shared forecasting systems. We show that over-adoption is guaranteed under strategic substitutes whenever adoption raises the dispersion of aggregate entry, holds in the large majority of admissible substitutes calibrations more generally and under moderate complementarity, and reverses only when agglomeration forces are sufficiently strong.

We also examine whether vendor competition mitigates the distortion. Multiple vendors reduce correlated errors only to the extent that their underlying data and models are independent. A descriptive case study of the restaurant site-selection analytics market suggests substantial overlap in upstream inputs and forecasting practices. Under empirically disciplined calibrations, the observed market retains roughly three-quarters of the distortion generated by a single shared platform. Vendor competition therefore provides limited diversification when providers draw on similar data sources and forecasting architectures.

2 Motivation: Fast Food Site Selection and AI Exposure

Large quick-service chains often make location decisions in geographic lockstep. A traveler who finds a McDonald's near a highway interchange often expects a Burger King nearby, and that expectation is frequently correct. This section documents that regularity in county-level location data and asks whether the two chains select similar types of markets.

The exercise is descriptive rather than causal. Its purpose is to show that the object formalized in the model, cross-firm alignment in market selection, is visible in a familiar retail setting before any structure is imposed. The analysis does not attempt to distinguish among common demand fundamentals, agglomeration economies, imitation, franchising constraints, real-estate constraints, or shared analytics. Instead, it documents a simple empirical regularity. McDonald's and Burger King exhibit highly similar mappings from observable market characteristics into store counts. This similarity is especially pronounced in counties with greater exposure to AI-supported analysis.

2.1 Industry Setting and Data

We assemble a county-level snapshot of McDonald's and Burger King locations in the United States.[3] Brand locations come from OpenStreetMap, and each location is assigned to a county through point-in-polygon matching against Census county boundaries. We merge these counts with county characteristics from the American Community Survey, the Census Gazetteer, and the USDA Rural-Urban Continuum Codes, as well as the county-level AI Geographic Exposure (AIGE) index of Felten et al. (2021).

The AIGE index measures how exposed a county's employment structure is to advances in artificial intelligence. Following Felten et al. (2021), the measure is constructed by combining occupation-level AI exposure scores with local employment composition. Counties with higher AIGE therefore have industry and occupation mixes that are more exposed to AI-related technolo-

[3] We focus on McDonald's and Burger King because they are large national chains in the same product segment and have broad geographic footprints. In the 2025 QSR 50, which ranks U.S. limited-service chains by 2024 system-wide sales, McDonald's is the largest chain overall and the largest burger chain, while Burger King ranks eighth overall and third among burger chains. Their common segment helps hold product category, service format, and target occasion relatively fixed, while their national footprints make county-level variation informative about broad site-selection patterns rather than the expansion strategy of a narrow regional brand. Both chains maintain substantial in-house real-estate and analytics capabilities, making this pair, if anything, a conservative setting for the shared-platform channel. Appendix F identifies chains that are clients of the same analytics vendors. Comparing alignment among disclosed co-clients against alignment among non-co-client pairs is a promising direction for future work with store-census data.

gies. We use AIGE as a stratifying variable rather than as a direct measure of platform adoption.

Two features of the data are important for interpretation. First, the brand layer is a partial collection. A zero outside the subset of states for which both brands were collected reflects absence of collection rather than absence of stores. We therefore restrict attention to counties in states where both brands were collected and treat a zero within that subset as a genuine absence. Second, OpenStreetMap is volunteered geographic information rather than a store census, so some locations may be present but untagged. The estimates should therefore be interpreted as descriptive rather than exhaustive. After restricting to counties with both brands collected and non-missing covariates, the sample contains 1,987 counties.

2.2 Empirical Approach

The empirical exercise documents alignment in market selection rather than a causal effect of AI exposure on store entry. A direct regression of store counts on AI Geographic Exposure is difficult to interpret because AIGE is constructed from local industry and occupation composition, which is also related to demand, density, commuting patterns, and commercial development. We therefore use AIGE only as a stratifying variable. The object of interest is whether McDonald's and Burger King map observable market characteristics into store counts in similar ways.

We estimate chain-specific relationships between store counts and county fundamentals and compare the resulting coefficient vectors. This asks whether the two chains place similar weights on the same observable market characteristics when evaluating markets. All continuous fundamentals are standardized before estimation so that coefficients are comparable across variables (Gelman 2008). We summarize cross-chain alignment using three measures: cosine similarity of the coefficient vectors, correlation of the coefficient vectors, and correlation between predicted market scores from the two chain-specific models. The first two measure similarity in estimated slopes. The third measures whether the two models rank counties similarly in predicted attractiveness, treating fitted values as informative descriptive objects even when the exercise is not causal (Mullainathan and Spiess 2017).

As a complementary diagnostic, we estimate a stacked chain-by-county specification with chain fixed effects and interactions between a Burger King indicator and the market fundamentals. This yields a Wald test of slope equality across chains (Paternoster et al. 1998). Exact equality is not the relevant benchmark. With a large county sample, economically small differences can gen-

erate statistical rejection even when the two chains rank markets similarly. We therefore view the Wald test as a diagnostic and place greater weight on the directional evidence from the coefficient and score-based measures.

Finally, we split counties at the median of the AI Geographic Exposure index and recompute the alignment measures separately in low- and high-exposure subsamples. This asks whether market-selection alignment is more pronounced in counties whose employment structure is more exposed to AI. Consistent with the descriptive approach emphasized by Einav and Levin (2014), the goal is to document a robust empirical pattern that motivates the theoretical mechanism developed below.

2.3 Co-Location

The two brands are present in the same counties far more often than their separate frequencies would imply. Table 1 cross-tabulates indicators for at least one McDonald's location and at least one Burger King location. Among counties with a McDonald's, 67.4 percent also contain a Burger King; among counties without a McDonald's, only 6.0 percent contain a Burger King. The two presence indicators have a Pearson correlation of 0.51. Conditional on population, density, income, urbanicity, and region, a county with a McDonald's is 11.9 percentage points more likely to also contain a Burger King, with a robust standard error of 2.7 percentage points.

Geographic alignment is therefore a prominent feature of the data. The two chains frequently co-locate, and this pattern survives coarse controls for observable market characteristics. Such alignment is not unique to the modern analytics era. Toivanen and Waterson (2005) document strategic interaction and co-location among U.K. hamburger chains using data that predate contemporary forecasting platforms. Our mechanism therefore does not seek to explain all cross-firm alignment. Instead, it highlights an additional channel through which shared forecasting infrastructure may increase alignment beyond what would arise from common demand fundamentals, agglomeration economies, or strategic interaction alone.

A cross-sectional snapshot cannot isolate this channel. The evidence here should therefore be interpreted as descriptive rather than causal. The purpose of the exercise is to document the object studied in the model, cross-firm alignment in market selection, and to show that it is economically important in a familiar retail setting.

Table 1: McDonald's and Burger King at the county level

	Burger King		
McDonald's	No	Yes	Total
No	408	26	434
Yes	506	1,047	1,553
Total	914	1,073	1,987

Notes: County-level observations in states for which both brands were collected in OpenStreetMap, with non-missing fundamentals ($N = 1{,}987$). Each cell is a count of counties. "Present" denotes at least one location of the brand in the county. Among McDonald's counties, the share with a Burger King is $1{,}047/1{,}553 = 0.674$; among non-McDonald's counties, it is $26/434 = 0.060$. OpenStreetMap is a partial, volunteered source and is not a store census.

2.4 Similar Market-Selection Rules

Co-location shows that the two chains often enter the same counties. The more relevant fact for the model is whether they select markets according to similar observable characteristics. Let n_{bi} denote the number of stores operated by brand $b \in \{\text{McDonald's}, \text{Burger King}\}$ in county i. We estimate brand-specific relationships between store counts and county fundamentals:

$$y_{bi} = \alpha_b + X_i'\beta_b + \varepsilon_{bi},$$

where y_{bi} is either $\log(1+n_{bi})$ or n_{bi} in a Poisson specification. The vector X_i includes population, population density, income, poverty, unemployment, education, vehicle access, Rural-Urban Continuum Code fixed effects, and Census-region fixed effects. Continuous fundamentals are standardized so that coefficients can be compared across brands.

The two chains differ in scale. McDonald's has 4.85 locations per county on average, compared with 2.11 for Burger King. Conditional on scale, however, the observable market-selection rules are nearly identical. In both the log-count and Poisson specifications, the coefficient vectors on standardized fundamentals are almost collinear, and the predicted market scores for the two chains are highly correlated. Observable fundamentals therefore rank counties in almost the same way for McDonald's and Burger King, even though McDonald's has the larger footprint.

A stacked chain-by-county specification provides a complementary test. We estimate a model with chain fixed effects and interactions between the Burger King indicator and the market fun-

Table 2: Similarity in market-selection rules

Specification	Cosine similarity	Coeff. correlation	RMS distance	Score correlation
$\log(1+\text{stores})$	0.9997	0.9997	0.0878	0.9964
Poisson count	0.9986	0.9985	0.0297	0.9885

Notes: The table compares chain-specific regressions of McDonald's and Burger King store counts on the same set of county fundamentals. Continuous fundamentals are standardized before estimation. Cosine similarity, coefficient correlation, and RMS distance are computed using the coefficient vectors on standardized continuous fundamentals. Score correlation is the correlation between fitted market scores from the two chain-specific models. Both specifications include Rural-Urban Continuum Code and Census-region fixed effects.

damentals. In the log-count and presence specifications, equality of slopes is rejected, largely because the two chains differ in scale and in their sensitivity to population. In the Poisson count specification, however, equality of the continuous-fundamental slopes is not rejected ($p = 0.238$), although equality is rejected once interactions with urbanicity and region are included. We read the Poisson result as the more informative of the two for the present question: the dependent variable is a store count with a substantial mass at zero, for which the Poisson conditional mean is the natural specification, whereas the $\log(1+n)$ transform compresses scale differences into the slope and rejects largely because the two chains differ in footprint size rather than in how they rank markets. The directional diagnostics below, which do not hinge on either functional form, make the same point more robustly.

This distinction is informative. Exact coefficient equality is too strong a benchmark for shared market selection. The relevant descriptive fact is directional alignment: the same observable fundamentals rank markets in nearly the same way for both chains. The two chains need not have identical coefficients for their entry rules to be highly aligned in the markets they identify as attractive.

2.5 AI Exposure as a Stratifying Variable

We do not use AI Geographic Exposure as a causal regressor in this exercise. A county's AI exposure reflects its industry and occupation mix, and is therefore related to the same market fundamentals that shape retail entry. We instead use AI exposure as a stratifying variable. Counties are split at the sample median of the AI Geographic Exposure index, and the market-selection similarity measures are recomputed separately for low- and high-exposure counties.

Table 3: Market-selection similarity by AI exposure

Sample	N	Cosine similarity	Coeff. correlation	RMS distance	Score correlation
Full sample	1,987	0.9997	0.9997	0.0878	0.9964
Low AIGE	993	0.9943	0.9941	0.0875	0.9849
High AIGE	994	0.9991	0.9991	0.1010	0.9978

Notes: The table reports similarity measures from chain-specific $\log(1+\text{stores})$ regressions estimated separately in the full sample, below-median-AIGE counties, and above-median-AIGE counties. AIGE is used only to split the sample, not as a treatment variable or direct regressor. The score correlation is the correlation between fitted market scores from the McDonald's and Burger King models.

The similarity is high in both subsamples, but stronger in high-AIGE counties. The predicted-score correlation is 0.9849 in low-AIGE counties and 0.9978 in high-AIGE counties, with coefficient-based measures showing the same pattern. This does not imply that AI exposure causes chains to behave more similarly. Rather, counties whose employment structures are more exposed to AI are also places where the two chains' observable market-selection rules are especially closely aligned.

These patterns motivate the model in two ways. First, McDonald's and Burger King exhibit substantial geographic alignment. Second, they map observable market characteristics into store counts in highly similar ways, especially in high-AIGE counties. This alignment is consistent with several mechanisms, including common demand fundamentals, agglomeration economies, imitation, franchising constraints, and shared analytics. The model focuses on the last channel and asks what happens when firms evaluate markets through related forecasting systems. In that case, entry decisions may become correlated in ways that firms do not fully internalize.

The descriptive evidence is consistent with this mechanism but does not identify it. Similar market-selection rules can arise from common fundamentals, strategic interaction, or shared forecasting infrastructure. Distinguishing among these channels requires variation that shifts the forecasting process without changing underlying market conditions, an identifying challenge that we leave to future work. The purpose of the exercise is therefore not to test the mechanism directly, but to document the empirical object that the model seeks to explain.

As a supplementary diagnostic, we decompose the difference in predicted-score alignment between high- and low-AIGE counties across the seven continuous market fundamentals using exact Shapley attribution.[4] The resulting contributions are small, partially offsetting, and imprecisely

[4]Shapley-value attribution is widely used in machine-learning interpretability (Lundberg and Lee 2017). We apply the exact decomposition of Shapley et al. (1953) to the high-minus-low AIGE difference in predicted-score alignment.

estimated. No individual fundamental accounts for a statistically distinguishable share of the difference. We therefore interpret the decomposition only as a diagnostic. The stronger alignment observed in high-AIGE counties does not appear to be driven by any single observable characteristic, but instead reflects broad-based similarity across market fundamentals.

3 Model

The model studies retail entry when firms choose both whether to enter a market and whether to rely on proprietary research or a shared forecasting platform. The central tradeoff is simple. Shared platforms can improve forecast accuracy by reducing idiosyncratic noise, but they also route decisions through a common processing layer that may generate correlated errors across adopters. The model isolates this tradeoff in the simplest environment that permits heterogeneous adoption and strategic interaction.

A key distinction is between primitives and outcomes. The variance of the shared processing component is a primitive of the information structure, whereas platform adoption is an endogenous equilibrium outcome that reflects firms' responses to those informational conditions.[5]

3.1 Environment

A unit continuum of retail chains $f \in [0,1]$ decides whether to enter a local market. The entry decision is binary, $x_f \in \{0,1\}$, and profits are

$$\Pi_f = (\theta + \delta \bar{x} + \varepsilon_f) x_f,$$

where $\theta \sim \mathcal{N}(0, \sigma_\theta^2)$ is unobserved local profitability, $\bar{x} \equiv \int_0^1 x_g \, dg$ is the aggregate entry rate, and $\varepsilon_f \sim \text{Uniform}[-\bar{\varepsilon}, \bar{\varepsilon}]$ is an idiosyncratic cost shock. The parameter $\delta \in (-1, \infty)$ governs strategic interaction. Positive values capture agglomeration economies, such as food halls, outlet centers, or fast-food strips where rival entry expands demand; negative values capture business stealing in markets with limited demand. The lower bound $\delta > -1$ plays no role in equilibrium existence, which is governed by the upper bound on δ collected in the admissible region $\mathscr{D}$ of Section 4; it instead bounds the strength of business stealing so that the mean entry rate $K_0 = 1/(2(1-\delta))$

[5] Throughout the analysis, σ_ξ^2 is treated as an exogenous property of the forecasting technology.

stays above $1/4$, keeping the interior regime the relevant benchmark.

The uniform cost shock is a payoff-smoothing device in the spirit of the global-games literature (Carlsson and van Damme 1993, Morris and Shin 1998); its role here is not equilibrium selection but linearity, as it makes entry probabilities linear in posterior beliefs. We work in the regime where this linear approximation is informative.[6] We normalize $\bar{\varepsilon} = 1/2$, so entry probability responds one-for-one to the posterior mean. This normalization is without loss of generality and does not affect the qualitative results.

3.2 Information and Processing Technologies

The model distinguishes between two approaches to forecasting. Firms can either generate forecasts internally using proprietary research or rely on a shared external platform. This distinction has become increasingly important as firms adopt AI-based forecasting technologies. Vendor-provided systems can improve forecasting performance by pooling data and analytical resources across users, but they may also reduce differentiation when many firms rely on similar tools.

This make-versus-buy margin has received growing attention in the empirical literature on firm AI. Babina (2026) identify it as a key dimension of AI adoption, while related evidence shows that firms investing more heavily in AI exhibit greater systematic risk (Babina et al. 2024) and that firms relying on external providers behave differently from those developing capabilities internally (Calvino and Fontanelli 2026). Our focus is on a question that this literature has largely left open: what are the welfare consequences when many firms rely on the same forecasting infrastructure?

The signal structure below formalizes this margin. Custom research generates idiosyncratic forecast errors that differ across firms. Platform-based forecasting may reduce this idiosyncratic variation, but it also introduces a common processing component shared across adopters. A chain using custom research observes

$$\tilde{s}_f^C = \theta + u_f^C,$$

[6] Formally, linearity holds whenever $|\mathbb{E}_f[\theta] + \delta \mathbb{E}_f[\bar{x}]| \leq \bar{\varepsilon}$ pointwise; a sufficient condition is $\bar{\varepsilon} > 4\sqrt{\mathrm{Var}(\theta) + \mathrm{Var}(\bar{x})}$. Some calibrations used in the figures lie outside this bound to illustrate qualitative regimes. The closed-form expressions of Proposition 1 remain well-defined wherever $\Delta(q) > 0$, and the qualitative comparative statics in Theorem 1 extend to a truncated-equilibrium model in which entry probabilities are clipped to $[0,1]$. Appendix D computes this clipped-rule equilibrium directly under a saturating logistic entry probability: the $\mathscr{W}_4$ wedge is exactly invariant to the replacement because it contains no equilibrium object, while the $\mathscr{W}_3$ wedge retains its sign and its monotonicity in q, with a rule-sensitive magnitude that converges to the linear closed form in the interior-regime limit.

where $u_f^C \sim \mathcal{N}(0, \sigma_C^2)$ is idiosyncratic noise. A chain using the shared platform observes

$$\tilde{s}_f^A = \theta + \xi + u_f^A,$$

where $u_f^A \sim \mathcal{N}(0, \sigma_A^2)$ is idiosyncratic noise and $\xi \sim \mathcal{N}(0, \sigma_\xi^2)$ is a common processing shock shared across all platform adopters. All shocks are mutually independent and independent of θ.

The distinction between the two technologies is therefore not only a matter of precision but of structure. Both signals are noisy measures of θ, and platform adoption may reduce idiosyncratic variance, so it is natural to expect $\sigma_A^2 < \sigma_C^2$ when data are pooled across users. At the same time, the platform replaces part of the idiosyncratic error with shared noise. This trade can be privately attractive while generating aggregate effects that no individual chain internalizes.

The shared component ξ represents forecast errors that originate in the processing technology rather than in market fundamentals. The motivation comes from a recurring lesson in computer science: prediction systems can exhibit similar measured accuracy while sharing latent sources of failure. Different pipelines may perform similarly in validation exercises yet make correlated mistakes in deployment because they rely on similar data, modeling choices, or software infrastructure (D'Amour et al. 2022, Sculley et al. 2015). When many firms rely on related forecasting systems, these mistakes can become common across adopters even when each system appears accurate in isolation.

This structure differs from the public-signal benchmark of Morris and Shin (2002). There, a common signal induces correlation because it is informative about the underlying state. As more agents observe the signal, actions become more aligned, but inference also improves. Here the shared component carries no information about θ. It survives aggregation as a common error, so joint exposure does not refine beliefs and instead transmits the same mistake across adopters.[7]

One way to interpret ξ is as a latent processing factor. When many chains rely on the same forecasting platform, it is as if they employ the same analyst, whose biases, training, and blind spots affect all users simultaneously. The variance σ_ξ^2 measures the importance of this shared component. The assumption that ξ is orthogonal to θ captures that these errors originate in the forecasting process rather than in the underlying market.

[7] This differs from environments such as Bergemann et al. (2015), where common signals may remain informative even when they confound multiple sources of uncertainty.

3.3 Adoption Stage

Each chain draws a switching cost $c_f \sim F$ independently, where F is a continuous distribution on $\mathbb{R}$ with positive density on its support. The chain chooses its forecasting technology to maximize expected entry profits net of this cost, taking the adoption rate q as given. An adoption equilibrium is a rational-expectations fixed point in which individual choices reproduce the aggregate:

$$q^* = F(\Delta\Pi(q^*)),$$

where $\Delta\Pi(q) \equiv \pi^A(q) - \pi^C(q)$ is the private gain from adoption. If F has bounded support $[0, \bar{c}]$, corner outcomes arise when $\Delta\Pi(0) \leq 0$ or $\Delta\Pi(1) \geq \bar{c}$. With unbounded support, such as the logistic specification used below, these cases are replaced by interior equilibria close to but bounded away from 0 or 1. This distinction does not affect the qualitative results.

3.4 Bayesian Filtering

Chains form beliefs using Bayesian filtering given their information and common knowledge of the data-generating process, including the presence and variance of the shared shock ξ. Let $V_A \equiv \sigma_\theta^2 + \sigma_\xi^2 + \sigma_A^2$ and $V_C \equiv \sigma_\theta^2 + \sigma_C^2$ denote unconditional signal variances, and define

$$\alpha_\theta^A = \frac{\sigma_\theta^2}{V_A}, \qquad \alpha_\xi^A = \frac{\sigma_\xi^2}{V_A}, \qquad \alpha_\theta^C = \frac{\sigma_\theta^2}{V_C}.$$

Then $\mathbb{E}_f^A[\theta] = \alpha_\theta^A \tilde{s}_f^A$, $\mathbb{E}_f^A[\xi] = \alpha_\xi^A \tilde{s}_f^A$, and $\mathbb{E}_f^C[\theta] = \alpha_\theta^C \tilde{s}_f^C$. Custom-research chains do not observe ξ and assign it its prior mean. The platform signal's common component is summarized by

$$\chi_A \equiv \alpha_\theta^A + \alpha_\xi^A = \frac{\sigma_\theta^2 + \sigma_\xi^2}{V_A},$$

the share of signal variance attributable to components common across adopters. This object governs how strongly the shared shock enters equilibrium behavior.

The assumption that firms know σ_ξ^2 is both consequential and conservative. It corresponds to a setting in which adopters understand the residual error structure of the forecasting platform and correctly account for it when forming beliefs. Such firms recognize that platform forecasts

contain a shared component unrelated to local demand and optimally discount that component when combining the forecast with prior information.

Many real-world adopters are unlikely to satisfy this benchmark. A regional grocer purchasing a site-selection package, a franchisee relying on analytics supplied by a franchisor, or a tenant receiving recommendations through a brokerage relationship may observe only a point forecast and have little information about the uncertainty underlying it. In such settings, firms may place too much weight on the platform signal and therefore react too strongly to the shared component.

The rational-filtering benchmark therefore provides a lower bound on the distortion generated by shared processing. The model isolates the welfare consequences of the information structure itself, abstracting from additional distortions that arise when firms misperceive the precision of platform forecasts. Any departure from correct filtering that increases reliance on the shared component strengthens the correlation in entry decisions and amplifies the resulting externality.[8]

4 Equilibrium with Heterogeneous Adoption

This section characterizes equilibrium behavior for a fixed adoption rate q. Throughout, we restrict attention to the admissible region

$$\mathscr{D} \equiv \left\{ \delta \in \mathbb{R} : -1 < \delta < \min\left(1 + \frac{\sigma_C^2}{\sigma_\theta^2},\ 1 + \frac{\sigma_A^2}{\sigma_\theta^2 + \sigma_\xi^2} \right) \right\},$$

on which the equilibrium denominator $\Delta(q)$ in Proposition 1 is strictly positive for all $q \in [0,1]$.

4.1 Best Responses and the Linear Fixed Point

A chain enters when its posterior expected profit is nonnegative,

$$\mathbb{E}_f[\theta] + \delta \mathbb{E}_f[\bar{x}] + \varepsilon_f \geq 0.$$

[8] The same logic extends beyond the binary adoption decision studied here. If firms instead choose the intensity with which they rely on a shared platform, private incentives continue to place excessive weight on the shared component because firms do not internalize the correlation imposed on other adopters. The correlated-evaluation externality therefore reappears on the intensive margin.

With $\varepsilon_f \sim \text{Uniform}[-1/2, 1/2]$, the entry probability is

$$P_f(\text{enter} \mid s_f) = \tfrac{1}{2} + \mathbb{E}_f[\theta] + \delta\, \mathbb{E}_f[\bar{x}],$$

in the interior region. We look for an equilibrium in which the aggregate entry rate is linear in the underlying shocks:

$$\bar{x} = K_0 + K_\theta \theta + K_\xi \xi.$$

In a continuum, idiosyncratic noise averages out. If k^A and k^C denote the slopes of entry probabilities with respect to signals for platform adopters and custom chains, consistency requires

$$K_\theta(q) = q\,k^A + (1-q)\,k^C, \qquad K_\xi(q) = q\,k^A.$$

Substituting the conjectured aggregate into each type's decision rule gives a linear system in (k^A, k^C).

Proposition 1 (Equilibrium loadings)**.** *For $q \in (0,1)$, the unique linear equilibrium has*

$$k^A(q) = \frac{\alpha_\theta^A}{\Delta(q)}, \qquad k^C(q) = \frac{\alpha_\theta^C\,(1 - \delta\, q\, \alpha_\xi^A)}{\Delta(q)}, \tag{1}$$

intercept $K_0 = 1/(2(1-\delta))$, and equilibrium denominator

$$\Delta(q) \equiv (1 - \delta\, q\, \chi_A)\,(1 - \delta(1-q)\,\alpha_\theta^C) - \alpha_\theta^A\,\alpha_\theta^C\,\delta^2\, q(1-q), \tag{2}$$

where $\chi_A = \alpha_\theta^A + \alpha_\xi^A = (\sigma_\theta^2 + \sigma_\xi^2)/V_A$. The equilibrium exists and is unique on $\mathscr{D}$.[9]

The loadings differ because platform adopters respond directly to both θ and ξ, while custom chains respond only to θ. The term $-\delta q \alpha_\xi^A$ in k^C captures the indirect effect of the shared shock on custom chains through aggregate entry. Although custom chains do not observe ξ, they anticipate that adopters respond to it. They therefore offset this effect when $\delta < 0$ and reinforce it when $\delta > 0$,

[9] The two bounds in $\mathscr{D}$ ensure $\Delta(0) > 0$ and $\Delta(1) > 0$, respectively, and concavity of $\Delta(q)$ implies positivity throughout $[0,1]$. The lower bound $\delta > -1$ plays no role in existence of the linear equilibrium, which is governed by the upper bound in $\mathscr{D}$. For $\delta \geq 1$, this requires $\sigma_C^2/\sigma_\theta^2 > \delta - 1$ and $\sigma_A^2/(\sigma_\theta^2 + \sigma_\xi^2) > \delta - 1$. Some calibrations lie in regions where these bounds bind. The closed-form expressions extend analytically to such cases but should then be interpreted as approximations to the truncated equilibrium described in Section 3.

with the adjustment becoming stronger as adoption rises or as the shared component accounts for a larger share of signal variance.

The boundary cases recover familiar benchmarks. When $q = 0$, the platform is absent and the equilibrium reduces to the standard private-signal entry game, with $k^C(0) = \alpha_\theta^C/(1-\delta\alpha_\theta^C)$. When $q = 1$, all chains adopt the platform and $k^A(1) = \alpha_\theta^A/(1-\delta\chi_A)$, where the denominator reflects that the platform signal contains both informative and uninformative common components.

4.2 Aggregate Behavior

Substituting the equilibrium loadings into the aggregate yields

$$\bar{x} = K_0 + K_\theta(q)\,\theta + K_\xi(q)\,\xi.$$

Using $\mathbb{E}[\theta] = \mathbb{E}[\xi] = 0$, the deviation from the mean is

$$\bar{x} - \mathbb{E}[\bar{x}] = K_\theta(q)\,\theta + K_\xi(q)\,\xi,$$

so aggregate variance is

$$\mathrm{Var}(\bar{x}) = K_\theta(q)^2\sigma_\theta^2 + K_\xi(q)^2\sigma_\xi^2.$$

The intercept $K_0 = 1/(2(1-\delta))$ can exceed unity for δ close to one, in which case the linear expression for $\bar{x}$ should be read as the interior-regime approximation discussed in Section 3. None of the objects analyzed below is affected: the variance decomposition, the adoption gain $\Delta\Pi(q)$, and every welfare wedge depend only on deviations of $\bar{x}$ from its mean, so the intercept drops out of all subsequent expressions.

This decomposition separates aggregate variation into two components. The first reflects fundamental tracking: the aggregate loads on θ, and its contribution scales with σ_θ^2. The second reflects shared processing: the aggregate loads on ξ, and its contribution scales with σ_ξ^2. In a custom-only economy, $K_\xi(0) = 0$, so aggregate fluctuations reflect fundamentals alone. Once platform adoption is positive, however, aggregate outcomes also inherit variation from the shared processing component.

The distinction matters because the two sources of variation behave differently under aggregation. Idiosyncratic noise averages out in the continuum, but the shared shock survives. As a

result, correlation across firms is no longer tied exclusively to information about demand. Part of that correlation reflects a common error embedded in the forecasting technology itself. The term $K_\xi(q)^2\sigma_\xi^2$ captures this source of aggregate variation and serves as the foundation for the welfare distortions studied below.

5 Adoption Incentives and Misalignment

This section characterizes adoption incentives and compares equilibrium adoption to social optima under alternative welfare criteria. The analysis separates two forces. The first is a classical entry externality: because each atomistic chain ignores its effect on the aggregate, the private equilibrium already departs from the chain-welfare optimum whenever entry decisions are strategic, even with no shared error. The second is the externality of interest: the shared component imposes costs on parties outside the chain population that no chain internalizes. We characterize the classical wedge first, then show that the shared component opens a further, external wedge under which a complete-welfare planner prefers strictly less adoption than a chain-welfare planner, regardless of strategic interaction.

5.1 Private Gain to Adoption

Let $\pi^A(q)$ and $\pi^C(q)$ denote expected profits for platform adopters and custom-research chains at adoption rate q, after substituting equilibrium behavior into the profit function. The private gain from adoption is $\Delta\Pi(q) \equiv \pi^A(q) - \pi^C(q)$, and an adoption equilibrium satisfies

$$q^* = F(\Delta\Pi(q^*)).$$

The dependence of $\Delta\Pi(q)$ on q reflects both forecast precision and strategic alignment. The platform changes the precision of a chain's information and also changes how closely its entry decision tracks rivals. The relative importance of these effects is governed by δ.

Proposition 2 (Monotonicity of the adoption gain)**.** *The function $\Delta\Pi(q)$ is independent of q when $\delta = 0$. For $\delta \neq 0$, $\Delta\Pi(q)$ is strictly increasing in q when $\delta > 0$ and strictly decreasing in q when $\delta < 0$. The sign result is established in closed form for strategic interaction near zero (Appendix G) and holds throughout the admissible region $\mathscr{D}$.*

When $\delta = 0$, chains care only about matching θ, so adoption depends on the signal-precision tradeoff. When $\delta > 0$, adoption becomes more attractive as it spreads because platform adopters track the shared component that increasingly shapes aggregate entry. When $\delta < 0$, the same mechanism reverses: chains benefit from differentiation, so the correlation induced by platform adoption reduces its marginal value.

The result follows from the derivative of $\Delta\Pi(q)$. Differentiating the closed-form expression derived in Appendix G gives

$$\frac{d\Delta\Pi}{dq} = \delta \cdot \frac{2\,\sigma_\theta^4\,\sigma_\xi^2}{V_A^3 V_C^2} \cdot \Gamma(q;\sigma,\delta),$$

where Γ is strictly positive over the admissible region $\mathscr{D}$; positivity reduces to a sign condition on a compact set, verified over 200,000 parameter draws in the replication package. The sign of $d\Delta\Pi/dq$ is therefore determined by δ.

5.2 Externalities with Affected Parties

Identifying the welfare costs of shared processing requires identifying who bears them. Chain profits, captured by the utilitarian benchmark below, do not exhaust the relevant population. The shared shock ξ generates two externalities, each affecting a different group.

The first externality operates at the aggregate level. When many firms rely on the same forecasting platform, aggregate entry becomes more sensitive to the shared processing component. This creates costs for agents whose payoffs depend on fluctuations in commercial activity around fundamentals, including lenders exposed to correlated default risk, landlords facing synchronized vacancies, and municipalities planning around local economic conditions. The relevant welfare object is the platform-driven component of aggregate variance, $K_\xi(q)^2\sigma_\xi^2$, which increases with adoption and survives aggregation.

The second externality operates at the location level. When multiple firms evaluate the same location through a shared platform, their forecasts inherit the same processing error. Evaluation mistakes therefore become positively correlated across firms. As adoption rises, independent screening becomes less effective because more firms rely on the same underlying signal. The corresponding welfare object is the expected covariance of evaluation errors across firms. This term scales as $q^2\sigma_\xi^2/V_A$, which reflects both the probability that two firms adopt the platform and the importance

of the shared component in platform forecasts.

This second channel is related to the algorithmic-monoculture concerns studied by Kleinberg and Raghavan (2021) and Peng and Garg (2024). In those settings, the decided-about agents are individuals, and welfare depends on whether at least one decision-maker recognizes a favorable case. Here, the decided-about agent is a location, and the distortion arises because multiple firms evaluate the same market through correlated signals. The underlying concern is similar, but the relevant welfare object differs.

Remark 1 (A screening micro-environment for $\mathscr{W}_4$)**.** The baseline environment features a single market evaluated by the full continuum. The pairwise object in $\mathscr{W}_4$ corresponds to the following micro-environment. Suppose that, alongside the entry game, there is a unit continuum of ex-ante identical candidate locations, each screened by two chains drawn independently at random from the population, with each evaluator applying its equilibrium forecasting technology to a location-specific realization of (θ, ξ, u). For a given location, the normalized covariance of the two evaluators' forecast errors equals σ_ξ^2/V_A when both evaluators are platform adopters and zero otherwise, so the expected normalized error covariance across locations is exactly $q^2\,\sigma_\xi^2/V_A$, the penalty term in $\mathscr{W}_4$. The weight γ_4 prices the loss of an independent second opinion: when evaluation errors are uncorrelated, a location wrongly screened out by one chain retains a chance of being recognized by the other, whereas correlated errors eliminate this diversification, in the spirit of Kleinberg and Raghavan (2021). Because a location passed over by both evaluators in error and a location entered by both in error are both costly, these losses do not net out across locations even though the shared error is mean zero; the relevant social loss is therefore the covariance itself rather than its (zero) expectation. The pairwise structure is a normalization, not a substantive restriction: if instead each location is screened by $n \geq 2$ evaluators drawn independently from the population, the expected number of co-adopting evaluator pairs is $\binom{n}{2}q^2$, each contributing a normalized error covariance σ_ξ^2/V_A, so the penalty is $\binom{n}{2}\,q^2\,\sigma_\xi^2/V_A$. The quadratic dependence on q and the linear dependence on σ_ξ^2/V_A are therefore properties of the encounter structure for every n; only the multiplicative constant changes, and it is absorbed into γ_4. The δ-invariance and strict positivity of the external $\mathscr{W}_4$ wedge in Theorem 1 hold unchanged for any fixed n.

The two externalities also differ in their dependence on equilibrium behavior. The aggregate-volatility externality depends on equilibrium loadings, while the correlated-evaluation externality

depends only on the signal structure and adoption rate. Both originate from the orthogonality of ξ to θ, but they operate through distinct channels.

5.3 Welfare Specifications

We consider three welfare specifications. The first is a utilitarian benchmark that aggregates chain profits. The second and third correspond to the two externalities identified above. A fourth specification, isolating fundamental-tracking performance, is discussed in Remark 3.[10]

The utilitarian benchmark, denoted $\mathscr{W}_1$, is

$$\mathscr{W}_1(q) = q\,\pi^A(q) + (1-q)\,\pi^C(q).$$

This specification captures the welfare of chains alone. As Proposition 3 below shows, the private equilibrium does not in general maximize $\mathscr{W}_1$: a strategic externality separates the two whenever $\delta \neq 0$. This classical wedge is logically separate from the shared-processing externalities that are the focus of this paper, and the two add.

The aggregate-volatility specification, denoted $\mathscr{W}_3$, augments $\mathscr{W}_1$ with a penalty on the platform-driven component of aggregate-entry variance:

$$\mathscr{W}_3(q) = \mathscr{W}_1(q) - \gamma_3\,K_\xi(q)^2\,\sigma_\xi^2, \qquad \gamma_3 > 0.$$

The term $K_\xi(q)^2\sigma_\xi^2 = q^2(k^A(q))^2\sigma_\xi^2$ is the contribution of the shared shock to aggregate variance. It depends on equilibrium loadings and therefore on the strategic-interaction parameter δ. The parameter γ_3 measures the marginal social cost of platform-induced volatility.

The decided-about-agent specification, denoted $\mathscr{W}_4$, penalizes correlated evaluation at the level of individual locations:

$$\mathscr{W}_4(q) = \mathscr{W}_1(q) - \gamma_4\,q^2\,\frac{\sigma_\xi^2}{V_A}, \qquad \gamma_4 > 0.$$

The term $q^2\sigma_\xi^2/V_A$ equals the probability that two evaluators both adopt the platform times the share of an adopter's signal variance attributable to the shared component. The normalization by V_A reflects that a location screened by multiple adopters bears the shared component relative to

[10] The indices 1, 3, and 4 are used for the specifications analyzed in this section, with the index 2 reserved for the fundamental-tracking criterion $\mathscr{W}_2$ of Remark 3. This numbering matches the figures and the replication code.

total signal uncertainty, not its absolute variance alone. The parameter γ_4 measures the marginal cost of correlated evaluation.

The distinction between $\mathscr{W}_3$ and $\mathscr{W}_4$ is structural. The $\mathscr{W}_3$ criterion penalizes realized aggregate behavior and therefore depends on equilibrium responses. The $\mathscr{W}_4$ criterion instead penalizes correlation in signals and depends only on adoption and the underlying information structure. As a result, the external $\mathscr{W}_4$ wedge at any fixed adoption rate is independent of strategic interaction, whereas the external $\mathscr{W}_3$ wedge varies with δ through its effect on equilibrium behavior.[11]

The penalty weights γ_3 and γ_4 determine the scale of the corresponding welfare losses but are not separately identified by the model. The identified content of the analysis is therefore the sign of each external wedge and its variation across environments rather than the absolute magnitude of welfare losses.

A natural alternative is the fundamental-tracking criterion $\mathscr{W}_2$ studied in the social-value-of-information literature (Angeletos and Pavan 2007, Morris and Shin 2002). This criterion captures costs borne within the chain population because deviations from fundamental tracking ultimately affect firms' own profits. Unlike $\mathscr{W}_3$ and $\mathscr{W}_4$, however, $\mathscr{W}_2$ is not an additive perturbation of $\mathscr{W}_1$. The clean external-wedge decomposition of Theorem 1 therefore does not apply, and Remark 3 shows that the resulting wedge is sign-indeterminate.

The distortion studied here instead arises from costs borne outside the chain population. A planner concerned with both objectives would naturally combine the two criteria, in which case the external wedge identified below would operate alongside any distortion associated with imperfect fundamental tracking.

5.4 Marginal Chain Welfare and the Classical Entry Externality

One might expect the linear-Gaussian continuum to deliver *team efficiency*, with the private equilibrium maximizing $\mathscr{W}_1$. It does not. The atomistic chain takes the law of motion of the aggregate $\bar{x}$ as given and therefore does not internalize how the population's technology mix moves the equilibrium loadings, and with them the strategic component $\delta\,\mathbb{E}[x_f\,\bar{x}]$ of every chain's profit. The result is a strategic externality whenever $\delta \neq 0$.

[11] The two criteria are not nested. Their ratio is $(k^A(q))^2 V_A$, which depends on δ through equilibrium loadings, so no choice of (γ_3, γ_4) equates them across environments. More generally, any welfare criterion that depends only on adoption and the information structure, rather than on equilibrium responses, generates an external wedge whose sign and magnitude at fixed q are independent of strategic interaction.

Proposition 3 (Marginal chain welfare)**.** *For all $q \in (0,1)$,*

$$\frac{d\mathscr{W}_1}{dq} = \Delta\Pi(q) + \Psi(q), \qquad \Psi(q) \equiv \frac{\delta}{2}\frac{d}{dq}\mathrm{Var}(\bar{x}) = \frac{\delta}{2}\frac{d}{dq}\left[K_\theta(q)^2\sigma_\theta^2 + K_\xi(q)^2\sigma_\xi^2\right].$$

Hence $d\mathscr{W}_1/dq = \Delta\Pi(q)$ if and only if $\delta = 0$. For $\delta \neq 0$ the private equilibrium q^* does not coincide with the chain-welfare optimum q_1^{soc} defined by $F^{-1}(q_1^{soc}) = d\mathscr{W}_1/dq$: atomistic chains do not internalize their effect on the shared strategic surplus, exactly the entry externality of Mankiw and Whinston (1986) and Spence (1976). The sign of the resulting adoption distortion is $\mathrm{sgn}(\Psi(q)) = \mathrm{sgn}\big(\delta \cdot \frac{d}{dq}\mathrm{Var}(\bar{x})\big)$, which is calibration-dependent because $\frac{d}{dq}\mathrm{Var}(\bar{x})$ need not be signed uniformly over q.

The term $\Psi(q)$ has an envelope interpretation. Each chain's decision rule is a best response to the aggregate, so a marginal change in q affects a chain's expected profit only through the equilibrium response of $\bar{x}$: as adoption rises, the loadings (k^A, k^C) adjust, the dispersion of aggregate entry changes, and the strategic component $\delta\,\mathbb{E}[x_f\bar{x}]$ of every chain's profit moves with it. Summed across the population, these uninternalized responses equal exactly $\Psi(q)$. Because firms choose forecasting technologies on private profits alone, this feedback enters no chain's decision, and the chain-welfare planner and the decentralized market generally choose different adoption levels.

This is the familiar wedge from the free-entry literature, translated from the entry margin to the technology-adoption margin. It is present even when $\sigma_\xi^2 = 0$ and therefore exists independently of any shared-processing considerations. Appendix G shows that this force is summarized by the term Ψ in the planner's first-order condition.

The shared-processing externality operates in addition to this classical distortion. To isolate the cost borne outside the chain population, we compare a chain-welfare planner, who maximizes $\mathscr{W}_1$, with a complete-welfare planner, who additionally internalizes the external cost. The difference between the two objectives is the object of interest. Because $\mathscr{W}_1$ enters both planners' problems identically, it cancels in the comparison. The resulting external wedge therefore does not inherit the strategic term Ψ and does not depend on any team-efficiency argument.

Remark 2 (Why the external wedge is robust)**.** The external wedge in Theorem 1 is the difference between the marginal values of adoption under $\mathscr{W}_1$ and $\mathscr{W}_k$. Because the two welfare criteria differ only through the penalty term, the contribution of $\mathscr{W}_1$ cancels exactly. This cancellation is algebraic and does not rely on a continuum of firms, linear entry probabilities, or any team-efficiency result.

Consequently, finite-N and nonlinearity concerns affect the strategic term Ψ but not the external wedge itself. In particular, the $\mathscr{W}_4$ penalty depends only on the adoption rate and the information structure through $q^2\sigma_\xi^2/V_A$. Its contribution to the external wedge therefore survives in finite-N and nonlinear environments. Appendix D confirms this: the $\mathscr{X}_4$ wedge is exactly invariant to a saturating entry rule, while the $\mathscr{X}_3$ wedge preserves its sign and its monotonicity in q.

5.5 Misalignment Theorem

The central welfare result is that internalizing either external cost calls for strictly less adoption. The relevant object is the wedge between the chain-welfare planner and the complete-welfare planner, which isolates the part of the distortion attributable to shared processing.

Theorem 1 (Robust external wedge)**.** *Fix $\delta \in \mathscr{D}$ and $\sigma_\theta^2, \sigma_A^2, \sigma_C^2, \sigma_\xi^2 > 0$. Define the external wedge under criterion $k \in \{3,4\}$ as $\mathscr{X}_k(q) \equiv d\mathscr{W}_1/dq - d\mathscr{W}_k/dq$.*

(i) Under aggregate-volatility welfare $\mathscr{W}_3$,

$$\mathscr{X}_3(q) = \gamma_3 \frac{d}{dq}\left[K_\xi(q)^2\,\sigma_\xi^2\right] > 0 \quad \textit{for all } q \in (0,1].$$

The wedge sign is invariant in δ, while its magnitude varies with δ through equilibrium loadings.

(ii) Under decided-about-agent welfare $\mathscr{W}_4$,

$$\mathscr{X}_4(q) = \gamma_4 \frac{d}{dq}\left[q^2\,\frac{\sigma_\xi^2}{V_A}\right] = \frac{2\,\gamma_4\, q\,\sigma_\xi^2}{V_A} > 0 \quad \textit{for all } q \in (0,1].$$

The wedge depends only on adoption and signal-structure primitives and is invariant in δ at any fixed q.

Consequently the complete-welfare optimum lies strictly below the chain-welfare optimum: $q_k^{soc} < q_1^{soc}$ for $k \in \{3,4\}$. Relative to the private equilibrium, the total gap decomposes as

$$\Delta\Pi(q) - \frac{d\mathscr{W}_k}{dq} \;=\; \mathscr{X}_k(q) \;-\; \Psi(q),$$

so the market over-adopts ($q^* > q_k^{soc}$) whenever the external wedge $\mathscr{X}_k$ dominates the strategic term Ψ. Whenever $\Psi(q) \le 0$, which under strategic substitutes holds exactly when $\frac{d}{dq}\mathrm{Var}(\bar{x}) \ge 0$,

the market over-adopts for any $\gamma_k > 0$; the comparison reverses only when the classical term is positive and exceeds the external wedge ($\Psi > \mathscr{X}_k$), which by Proposition 3 requires δ and $\frac{d}{dq}\mathrm{Var}(\bar{x})$ to share a sign with Ψ large in magnitude, as under strong agglomeration.

The two criteria share a common structure: the external wedge equals the marginal external cost of adoption. Under $\mathscr{W}_3$, this cost is the derivative of $K_\xi(q)^2\sigma_\xi^2$, which combines the primitive variance with the equilibrium amplification term $K_\xi(q)^2 = q^2(k^A(q))^2$. Its sign is invariant, but its magnitude varies with δ through the equilibrium loadings. Under $\mathscr{W}_4$, the relevant object is the derivative of $q^2\sigma_\xi^2/V_A$, which equals the encounter probability times the within-adopter variance share. Because equilibrium loadings do not appear, the wedge at a fixed adoption rate is independent of δ. The external result therefore does not depend on whether entry decisions are strategic complements or substitutes. It follows from two features of the environment: the presence of a common component unrelated to fundamentals and the fact that platform adoption increases exposure to that component.

The decomposition $\Delta\Pi - d\mathscr{W}_k/dq = \mathscr{X}_k - \Psi$ makes the two forces explicit. The shared-processing externality $\mathscr{X}_k$ always points toward less adoption. The classical strategic externality Ψ points toward less adoption under business stealing and toward more adoption under agglomeration whenever adoption raises the dispersion of aggregate entry, with both directions reversing when it lowers it. Whether the decentralized market over- or under-provides adoption relative to the complete-welfare optimum depends on the balance of these two forces. The contribution of this paper is the first term: a distinct external channel that is strictly signed and, under $\mathscr{W}_4$, independent of the strategic environment.[12]

The three calibrations of Figure 1 illustrate the wedges under moderate complementarity ($\delta = 0.7$); the parameter sweeps in the replication package extend the evaluation across the admissible region, including strategic substitutes. In the strategic-substitutes region, $\Psi \leq 0$ whenever $\frac{d}{dq}\mathrm{Var}(\bar{x}) \geq 0$, in which case $\mathscr{X}_4 - \Psi \geq \mathscr{X}_4 > 0$ and over-adoption is guaranteed. Because the variance derivative is not uniformly signed, the guarantee is conditional: across 30,000 random admissible draws with $\delta \leq 0$ and $\gamma_4 = 1$, the inequality $\mathscr{X}_4 - \Psi > 0$ holds in 96.6 percent of draws,

[12]The wedge admits a direct Pigouvian interpretation, but it must internalize both forces. The tax $\tau_k(q) = \mathscr{X}_k(q) - \Psi(q) = \Delta\Pi(q) - d\mathscr{W}_k/dq$ on platform adoption aligns the decentralized condition with the complete planner's first-order condition $F^{-1}(q) = d\mathscr{W}_k/dq$. The component that corrects shared processing specifically is $\mathscr{X}_k(q)$; the component $-\Psi(q)$ corrects the classical entry externality, and its sign follows that of Ψ in Proposition 3: it is a subsidy under agglomeration in the case $\frac{d}{dq}\mathrm{Var}(\bar{x}) > 0$, where $\Psi > 0$.

and every failure occurs at a calibration in which adoption lowers the dispersion of aggregate entry, so that $\Psi > 0$.

Of the two criteria, $\mathscr{W}_4$ provides the cleaner benchmark. Its external wedge depends only on adoption and signal-structure primitives. It is also invariant to the entry-choice rule by construction: $\mathscr{W}_1$ cancels in the wedge and the penalty $q^2\sigma_\xi^2/V_A$ depends only on adoption and the signal structure, so replacing the linear rule with any alternative leaves $\mathscr{X}_4$ unchanged exactly (Appendix D). The aggregate-volatility criterion $\mathscr{W}_3$ yields the same qualitative conclusion, but its magnitude inherits the equilibrium loadings and is therefore sensitive to the curvature of the choice rule, preserving its sign and its monotonicity in q under a saturating logistic rule while attenuating in level (Appendix D). We therefore view $\mathscr{W}_4$ as establishing the external-wedge result and $\mathscr{W}_3$ as corroborating it.

This invariance has an additional empirical advantage. The strategic parameter δ is difficult to identify because agglomeration and business-stealing effects often coexist and vary across settings. Moreover, δ enters both the classical wedge Ψ and the magnitude of $\mathscr{X}_3$. Under $\mathscr{W}_4$, no such identification is required. The external wedge depends only on adoption and the signal structure.

The decided-about-market penalty in $\mathscr{W}_4$ enters as an additive perturbation of chain welfare, with the encounter structure underlying it sketched in Remark 1. Two questions about this construction bear on the interpretation of the external wedge. First, is the screening loss a product of the entry environment, or an object appended to firm profits to obtain the result? Second, does the δ-invariance of $\mathscr{X}_4$ in Theorem 1(ii) reflect a property of the externality, or an artifact of evaluating it outside the entry game? Appendix E addresses both. It embeds the screening problem in a unified environment of markets each screened by $n \geq 2$ firms drawn from the population, derives the exposure $\binom{n}{2} q^2\sigma_\xi^2/V_A$ as an exact second moment of that environment rather than a posited penalty, and promotes the wedge to Proposition 9. The same analysis shows that the δ-invariance is diagnostic rather than incidental. The screening externality runs through the residual covariance of signals, which contains no equilibrium loading; strategic interaction would re-enter only through cross-market linkages, and then through the same channel as the aggregate-volatility wedge $\mathscr{X}_3$. The decided-about-market externality is therefore a screening externality, distinct from the coordination externality embedded in $\mathscr{X}_3$, and the distinction is precisely the dependence on δ.

Remark 3 (Misalignment under fundamental-tracking welfare $\mathscr{W}_2$ is generically indeterminate)**.**

A fourth specification isolates forecast accuracy at the chain level, treating coordination payoffs as transfers and penalizing deviations from the fundamental in the spirit of Angeletos and Pavan (2007):

$$\mathscr{W}_2(q) = -q\,(k^A(q)-1)^2 - (1-q)\,(k^C(q)-1)^2.$$

This criterion penalizes deviations of each type's signal loading from one, the loading at which entry passes signals through fully; an alternative tracking criterion that penalizes the mean-squared gap between aggregate entry and the fundamental, $\mathbb{E}[(\bar{x}-\mathbb{E}[\bar{x}]-\theta)^2]$, delivers the same qualitative conclusion, because both criteria measure objects internal to the chain population and both generate indirect terms in $(k^A)'(q)$ and $(k^C)'(q)$ whose signs are calibration-dependent. Unlike $\mathscr{W}_3$ and $\mathscr{W}_4$, $\mathscr{W}_2$ is not an additive perturbation of $\mathscr{W}_1$, so its wedge does not reduce to a clean penalty derivative and instead carries indirect loading terms.

The resulting wedge $d\mathscr{W}_1/dq - d\mathscr{W}_2/dq$ includes indirect terms in $(k^A)'(q)$ and $(k^C)'(q)$, so its sign is calibration-dependent. This specification therefore does not deliver a determinate prediction for over- or under-adoption.

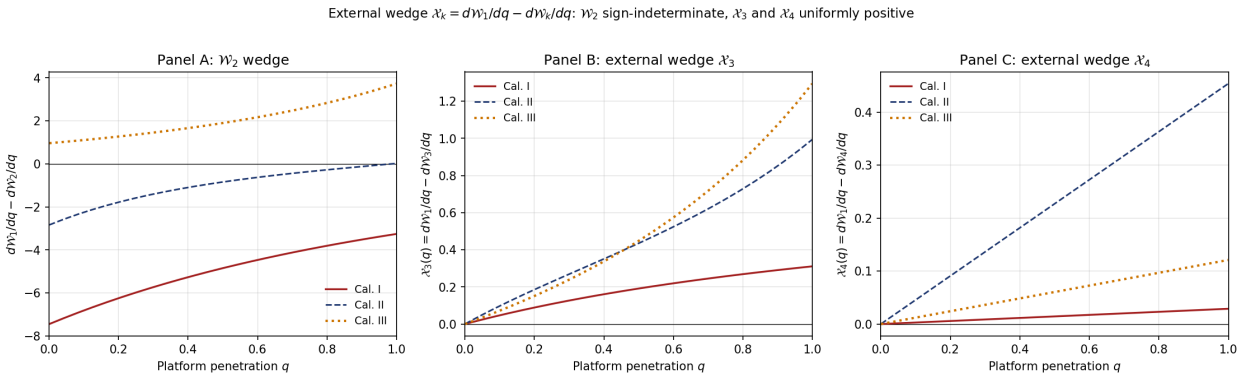
The contrast isolates the role of externalities. Under $\mathscr{W}_2$, costs are internal to the chain population. A determinate external wedge arises only when the planner accounts for costs borne outside the chain population, as in $\mathscr{W}_3$ and $\mathscr{W}_4$.

The contrast across the three panels of Figure 1 isolates the role of the welfare specification. Panel A shows that not every reasonable criterion delivers a determinate sign: $\mathscr{W}_2$ measures an object internal to the chain population, and the wedge is calibration-dependent. Panels B and C show that criteria targeting external costs yield a uniformly positive external wedge. The result therefore reflects a structural feature of the model rather than a particular welfare choice: whenever the planner internalizes costs borne outside the chain population, the optimal adoption level falls.

5.6 Residual Error versus Informative Commonality

The baseline model assumes $\xi \perp \theta$ to isolate the correlated-evaluation channel from the social value of public information. A natural question is whether the external wedge depends on the common component being entirely unrelated to fundamentals. To draw this boundary, suppose the

Figure 1: The external wedge under three welfare specifications.



Notes: Each panel plots a marginal welfare wedge as a function of platform penetration q, for three calibrations sharing $\delta = 0.7$ and $\gamma_k = 1$ but differing in the relative magnitudes of σ_A^2, σ_C^2, and σ_ξ^2. Calibration I: $(\sigma_\theta^2, \sigma_\xi^2, \sigma_A^2, \sigma_C^2) = (2.79, 0.05, 0.62, 0.27)$; Calibration II: $(1.23, 0.62, 0.88, 0.47)$; Calibration III: $(2.48, 0.21, 0.78, 1.46)$. Panel A shows the $\mathscr{W}_2$ wedge $d\mathscr{W}_1/dq - d\mathscr{W}_2/dq$: it is negative throughout for Calibrations I and II and positive throughout for Calibration III, with no interior zero-crossing, so the same primitives deliver no determinate prediction about over- or under-adoption relative to the $\mathscr{W}_2$ planner (Remark 3). Panel B shows the external wedge $\mathscr{X}_3(q) = d\mathscr{W}_1/dq - d\mathscr{W}_3/dq$ under aggregate-volatility welfare: it is strictly positive across all three calibrations and all $q \in (0,1]$, with magnitude depending on δ through equilibrium loadings (Theorem 1, part i). Panel C shows the external wedge $\mathscr{X}_4(q)$ under decided-about-agent welfare: it is strictly positive, linear in q, and invariant in δ (Theorem 1, part ii). Levels are in units of γ_k. Panels B and C jointly establish that the external wedge is robust across two structurally distinct externality penalties. The net private-vs-social comparison additionally subtracts the classical strategic term $\Psi(q)$ of Proposition 3, which is not plotted.

shared component can be decomposed as

$$\xi = \ell\theta + \zeta, \qquad \zeta \perp \theta, \qquad \mathrm{Var}(\zeta) = \sigma_\zeta^2,$$

where $\ell \in [0,1)$ is the loading on fundamentals. The platform signal becomes

$$\tilde{s}_f^A = (1+\ell)\theta + \zeta + u_f^A,$$

with total variance

$$V_A(\ell) \equiv (1+\ell)^2\sigma_\theta^2 + \sigma_\zeta^2 + \sigma_A^2.$$

The informative component changes signal precision. The residual component ζ is the part that remains after conditioning on fundamentals and is therefore the part to which the correlated-evaluation externality attaches.

Proposition 4 (Residual error versus informative commonality)**.** *Suppose the shared component*

admits the decomposition $\xi = \ell\theta + \zeta$*, where* $\zeta \perp \theta$*,* $\mathrm{Var}(\zeta) = \sigma_\zeta^2$*, and* $\ell \in [0,1)$*. Let*

$$\mathscr{D}(\ell) \equiv \left\{ \delta \in \mathbb{R} : -1 < \delta < \min\left(1 + \frac{\sigma_C^2}{\sigma_\theta^2},\ 1 + \frac{\sigma_A^2}{(1+\ell)^2\sigma_\theta^2}\right) \right\}$$

denote the admissible region under the transformed signal, and let $V_A(\ell) \equiv (1+\ell)^2\sigma_\theta^2 + \sigma_\zeta^2 + \sigma_A^2$*. Under the decided-about-agent welfare criterion* $\mathscr{W}_4$*, the external wedge is*

$$\mathscr{X}_4(q) \;=\; \frac{d\mathscr{W}_1}{dq} - \frac{d\mathscr{W}_4}{dq} \;=\; \frac{2\gamma_4 q\,\sigma_\zeta^2}{V_A(\ell)}.$$

Consequently:

(i) For all $q \in (0,1]$ *and* $\delta \in \mathscr{D}(\ell)$*, the wedge is strictly positive whenever* $\sigma_\zeta^2 > 0$ *and zero when* $\sigma_\zeta^2 = 0$*.*

(ii) Holding residual variance σ_ζ^2 *fixed, a larger informative component* ℓ *lowers the wedge only by increasing total signal variance* $V_A(\ell)$*, and the wedge vanishes only as* $\sigma_\zeta^2 \to 0$*.*

(iii) At any fixed adoption rate, the wedge is independent of δ*.*

The proposition separates the mechanism from standard public-information effects. In public-information models (Morris and Shin 2002, Angeletos and Pavan 2007, Hellwig and Veldkamp 2009), the common signal is informative about fundamentals, so its welfare effect depends on the tradeoff between coordination benefits and reduced independent learning. Here the external wedge is tied to the residual error ζ, which carries no information about θ. A larger informative component dilutes the residual share, but the wedge remains positive whenever residual error remains. The orthogonal baseline is the special case $\ell = 0$, in which σ_ξ^2 itself is the relevant primitive.

Remark 4 (Information-structure robustness)**.** Proposition 4 shows that the external wedge is robust to enriching the shared component with an informative part. The same logic extends to the comparison of information structures more broadly: what governs the externality is the residual common variance σ_ζ^2 that survives conditioning on fundamentals, not the joint informativeness of the signal. This is the sense in which the external-wedge result is a prediction about the information structure rather than about a particular parameterization, and it is robust across the structures a platform might plausibly implement (Bergemann and Morris 2013).

The policy implication follows directly. Improving common-signal informativeness can reduce the wedge by increasing $V_A(\ell)$, but it does not internalize the correlated-mistake channel. Mitigation requires reducing σ_ξ^2 itself, either by increasing structural independence across vendors, limiting overlap in upstream inputs, or auditing the processing layer that converts data into forecasts.

6 Four Regimes of Adoption

Adoption is determined by the fixed point $q^* = F(\Delta\Pi(q^*))$. The number and stability of equilibria depend on the shape of $\Delta\Pi(q)$ and on its boundary values. The first boundary admits a simple characterization.

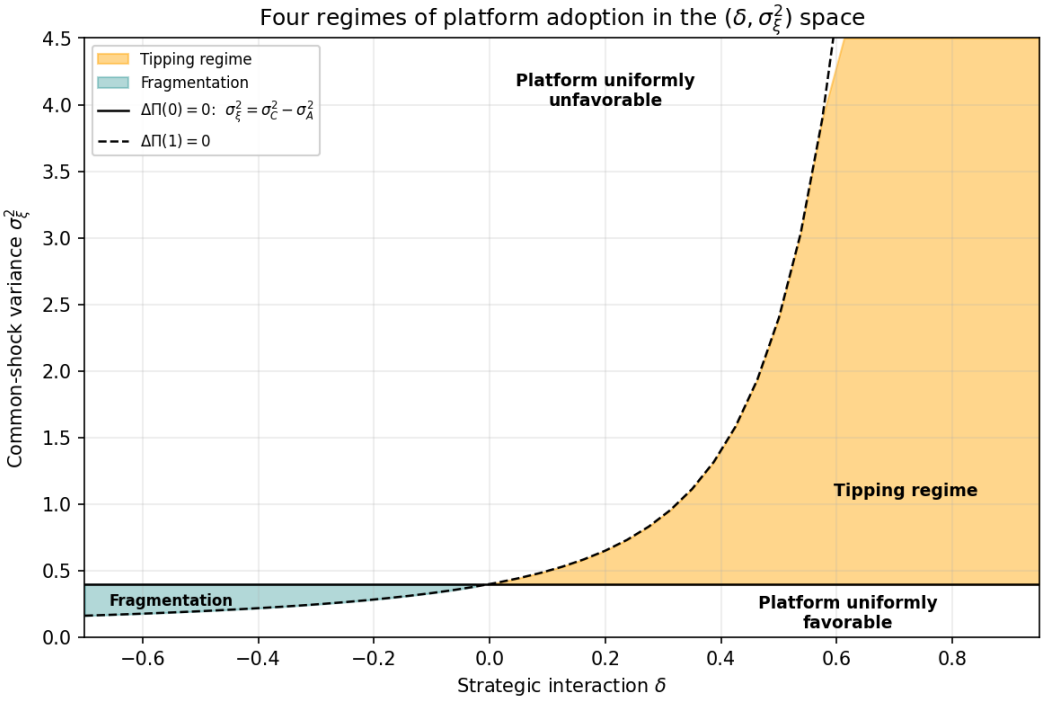
Lemma 1 (δ-independent threshold)**.** *The condition* $\Delta\Pi(0) = 0$ *holds if and only if* $\sigma_\xi^2 = \sigma_C^2 - \sigma_A^2$, *independent of* δ.

At $q = 0$, a marginal adopter faces the same aggregate as a non-adopter, so the coordination channel does not differentiate the technologies. Adoption depends only on signal quality. When $V_A = V_C$, platform and custom research deliver identical precision, and expected profits coincide. The platform is attractive in isolation if and only if its reduction in idiosyncratic noise, $\sigma_C^2 - \sigma_A^2$, exceeds the variance of the shared component.

The second boundary, defined by $\Delta\Pi(1) = 0$, depends on δ. As coordination motives strengthen, deviation from the platform becomes more costly, so chains tolerate a larger shared component before switching away. The corresponding threshold in σ_ξ^2 is increasing in δ. Numerical evaluation shows that this locus is monotone and convex and intersects the $q = 0$ boundary at $\delta = 0$.

Figure 2 plots these two loci in (δ, σ_ξ^2) space and partitions the parameter region into four regimes. While equilibrium configurations vary across regions, the external wedge of Theorem 1 is strictly positive throughout: strategic interaction shapes how adoption is sustained and whether the *net* private distortion is over- or under-adoption, but the component attributable to shared processing always points toward less adoption. The regimes therefore differ in the level and structure of adoption, not in the sign of the external wedge. Two cases are particularly informative: the tipping regime under agglomeration and the fragmentation regime under business stealing.

Figure 2: Four regimes of platform adoption in the (δ, σ_ξ^2) parameter space.



Notes: The phase diagram shows four adoption regimes as functions of the strategic-interaction parameter δ and the common-shock variance σ_ξ^2. The solid line is the locus $\Delta\Pi(0) = 0$, where the marginal platform deviator at zero penetration is just indifferent; this boundary is independent of δ (Lemma 1). The dashed curve is the locus $\Delta\Pi(1) = 0$, where the marginal custom-research deviator at full penetration is just indifferent; this curve is monotone increasing in δ and passes through $\sigma_\xi^2 = \sigma_C^2 - \sigma_A^2$ at $\delta = 0$, and is shown only over the range of δ on which the linear equilibrium is admissible. The shaded amber region is the tipping regime: a high-adoption fixed point coexists with a low-adoption fixed point, sustained by coordination motives even though the platform is intrinsically unfavorable in isolation. The shaded teal region is the fragmentation regime under business stealing, in which the unique adoption fixed point is interior. The horizontal axis is restricted to $\delta \in [-0.7, 0.95]$, inside the admissibility region $\mathscr{D}$ of Proposition 1 for this calibration; regions are not extrapolated beyond the plotted boundary loci. Calibration: $\sigma_\theta^2 = 1$, $\sigma_A^2 = 0.2$, $\sigma_C^2 = 0.6$.

6.1 Tipping Regime

When $\delta > 0$, $\Delta\Pi(0) < 0$, and $\Delta\Pi(1) > 0$, the fixed-point condition $q^* = F(\Delta\Pi(q^*))$ has at least one interior solution. If F is sufficiently curved relative to $\Delta\Pi$, so that $F'(\Delta\Pi(q))\Delta\Pi'(q) > 1$ on a non-degenerate interior interval, the equation admits three solutions: a stable low-adoption fixed point q_L^*, a stable high-adoption fixed point q_H^*, and an unstable threshold $q^\dagger$ separating their basins of attraction.[13] Figure 3 illustrates such a configuration for a logistic F with steepness parameter κ_F large enough for the curvature condition to hold.

[13] Stability refers to best-response iteration on the adoption mapping. A fixed point q^* is stable if $|F'(\Delta\Pi(q^*))\Delta\Pi'(q^*)| < 1$, where F' is evaluated at $\Delta\Pi(q^*)$ and $\Delta\Pi'$ at q^*.

In this region, platform adoption is privately unattractive in isolation but attractive once adoption is sufficiently widespread. At low penetration, exposure to the shared component is limited, and chains prefer custom research. At high penetration, aggregate entry is driven more strongly by the shared component, and chains benefit from aligning with rivals who respond to the same signal. Crossing $q^\dagger$ shifts the dominant force from precision to coordination, making adoption self-reinforcing.

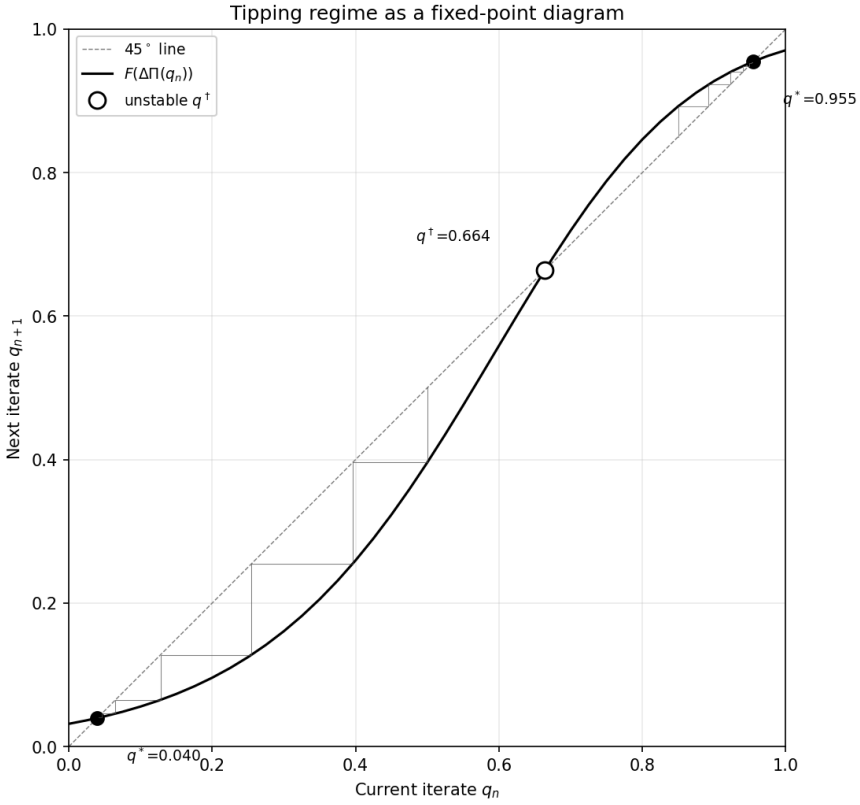
This structure parallels classic models of technology adoption under increasing returns. In Arthur (1989), small historical events are amplified through positive feedback, leading the market to lock in to a single technology. Here, the feedback operates through strategic interaction and shared processing rather than technological increasing returns alone. The inefficiency is also distinct: coordination may be privately valuable, but the shared component generates correlated errors that a complete-welfare planner would penalize.

This regime provides a precise formulation of shared platforms as coordination devices. The high-adoption fixed point is sustained even though adoption is privately unattractive in isolation, because no chain has an incentive to deviate once others adopt. The mechanism is strongest in sectors with agglomeration economies. In food halls, outlet centers, and clusters of competing fast-food restaurants, forecasting rival behavior can be as valuable as forecasting demand, so a platform that reveals common signals becomes valuable even if it is not informative about fundamentals.

Under either external-party welfare benchmark $\mathscr{W}_3$ or $\mathscr{W}_4$, a complete-welfare planner would locally prefer less adoption than the chain-welfare planner at each attracting fixed point, by the external wedge of Theorem 1. Whether each fixed point also lies above the complete planner's own optimum depends on the net sign $\mathscr{X}_k - \Psi$, which is positive whenever $\Psi \leq 0$ and can turn negative when $\Psi > 0$ is large, as under strong agglomeration with $\frac{d}{dq}\mathrm{Var}(\bar{x}) > 0$. The distinction between marginal and global comparisons is also important. The external wedge $\mathscr{X}_k > 0$ is a marginal statement at every interior q. The global comparison between q_L^* and q_H^* instead depends on welfare levels, $\mathscr{W}_k(q_L^*)$ versus $\mathscr{W}_k(q_H^*)$, and therefore on the relative magnitude of coordination gains and external costs.

Under $\mathscr{W}_3$, the externality scales with $K_\xi(q)^2$ and therefore inherits curvature from δ. For moderate δ, the low-adoption equilibrium is welfare-superior because the externality dominates the coordination gain, while for sufficiently large δ the ranking can reverse. Under $\mathscr{W}_4$, the externality scales with q^2 independently of δ, so the high-adoption equilibrium is more likely to be welfare-

Figure 3: The tipping regime as a fixed-point diagram.



Notes: The figure plots the next iterate $q_{n+1} = F(\Delta\Pi(q_n))$ against the current iterate q_n under best-response iteration on the static adoption map. The index n counts steps of this iteration, a graphical device used to locate the static fixed points of $q^* = F(\Delta\Pi(q^*))$ and assess their stability; it does not correspond to calendar time, which is reserved for the dynamic model of Section 8. Fixed points lie at intersections with the 45-degree line. Three intersections obtain: a stable low-adoption fixed point $q_L^* \approx 0.04$ and a stable high-adoption fixed point $q_H^* \approx 0.96$, separated by an unstable interior fixed point $q^\dagger \approx 0.66$ at which the slope of $F \circ \Delta\Pi$ exceeds unity. Best-response iteration converges to whichever stable fixed point lies in the same basin of attraction as the initial iterate: the wide left basin captures industries that begin from any moderate level of platform adoption, while only initial values above $q^\dagger$ tip the industry into the high-adoption fixed point. Calibration: $\delta = 0.7$, $\sigma_\theta^2 = 1$, $\sigma_\xi^2 = 0.6$, $\sigma_A^2 = 0.7$, $\sigma_C^2 = 0.85$, with logistic adoption distribution $F(x) = (1+e^{-\kappa_F x})^{-1}$ and $\kappa_F = 25$. This calibration sits inside the admissibility region $\mathscr{D}$ of Proposition 1, with $\Delta(q) > 0$ throughout $[0,1]$.

dominated. A full global ranking is beyond the scope of the present analysis; the external wedge remains positive at any interior equilibrium.

A regulator seeking to reduce synchronization must therefore consider both the external wedge and the ranking of equilibria. When the low-adoption equilibrium is preferred, policy must shift the economy across the unstable threshold. Measures such as disclosure requirements, diversification mandates, or support for alternative technologies operate by shrinking the basin of attraction of the high-adoption outcome and expanding that of the low-adoption equilibrium.

6.2 Adoption Fragmentation

When $\delta < 0$, $\Delta\Pi(0) > 0$, and $\Delta\Pi(1) < 0$, neither corner is a fixed point. Since $F \circ \Delta\Pi$ is continuous and strictly decreasing in q, the fixed-point equation admits a unique interior solution. The equilibrium features persistent heterogeneity, with a stable mix of adopters and non-adopters.

The mechanism reverses the coordination case. The platform improves precision and is therefore attractive in isolation, but it also increases correlation across chains. When entry decisions are strategic substitutes, this correlation is costly. As adoption rises, the marginal benefit declines, and equilibrium adoption settles where precision gains are offset by the cost of moving with rivals. Sectors where differentiation is central, such as specialty grocers, boutique retail, or niche restaurant concepts, should therefore exhibit partial adoption rather than convergence to a common platform.

The welfare implication is clean in this regime. By Theorem 1, the external wedge is positive; the fragmentation equilibrium over-adopts whenever the external wedge dominates the classical term Ψ, which is guaranteed whenever $\frac{d}{dq}\mathrm{Var}(\bar{x}) \geq 0$, so that $\Psi \leq 0$ by Proposition 3, and holds in the large majority of admissible substitutes calibrations (Section 5). Strategic substitution then limits the level of adoption but does not eliminate the externality. The tipping and fragmentation regimes therefore differ in the level of q^* and in whether the classical and shared-processing externalities reinforce or oppose, but the external wedge attributable to shared processing is excessive in both.

6.3 Boundary Regimes

The remaining cases are degenerate. When $\Delta\Pi(0) > 0$ and $\Delta\Pi(1) > 0$, the platform is privately attractive at all penetration levels and equilibrium adoption is bounded away from zero. When $\Delta\Pi(0) < 0$ and $\Delta\Pi(1) < 0$, the platform is unattractive throughout and adoption remains near zero. In both cases, strategic forces reinforce the underlying technological advantage or disadvantage rather than changing its direction.

7 Vendor Diversity and the Magnitude of Misalignment

The external-wedge result of Theorem 1 establishes that a complete-welfare planner prefers less adoption but does not determine the size of the wedge. Allowing multiple competing platforms introduces a second policy margin: the diversity of processing technologies. This margin matters because platforms may differ in interface or branding while relying on the same underlying data or models. The multi-platform setup below distinguishes substantive diversity from cosmetic diversity.

7.1 Multiple Processing Technologies

Suppose adopters partition across $J \geq 1$ platforms. Platform j generates a processing shock $\xi_j \sim \mathcal{N}(0, \sigma_\xi^2)$ and idiosyncratic noise $u_f^{A,j}$ with variance σ_A^2, so that an adopter of platform j observes

$$\tilde{s}_f^{A,j} = \theta + \xi_j + u_f^{A,j}.$$

Let q_j denote the share of chains using platform j, with $\sum_{j=1}^{J} q_j = q$. The shocks $\{\xi_j\}$ are jointly Gaussian with $\mathrm{Var}(\xi_j) = \sigma_\xi^2$ and pairwise correlation $\mathrm{Corr}(\xi_j, \xi_{j'}) = \rho \in [0,1]$ for $j \neq j'$. The single-platform model corresponds to $J = 1$.

The parameter ρ captures shared processing across platforms. When $\rho = 0$, platforms are structurally independent, with shocks arising from distinct models or data sources. When $\rho = 1$, platforms share a common core, so that $\xi_j = \xi_{j'}$ almost surely. Intermediate values capture partial overlap in inputs or methods.[14]

For tractability, vendor shares $\{q_j\}$ are treated as exogenous. Section 6 determines total adoption q endogenously; here we hold q fixed and isolate how its allocation across platforms affects welfare.

[14] Such overlap need not be confined to data inputs. DosSantos DiSorbo et al. (2026) document that frontier language models from every major provider exhibit the same systematic rigidity in handling policy exceptions, refusing warranted deviations that humans grant; because this bias is a property of how the underlying models are trained rather than of any one vendor's interface, downstream tools built on a shared foundation-model layer inherit correlated decision errors even when separately branded. This is a behavioral instance of the high-ρ regime.

7.2 Effective Concentration Index

Each adopter continues to observe a single processing shock, so the equilibrium loadings (k^A, k^C) are unchanged from Proposition 1. The difference arises in aggregation:

$$\bar{x} = \mathbb{E}[\bar{x}] + K_\theta(q)\,\theta + \sum_{j=1}^{J} K_{\xi,j}\,\xi_j, \qquad K_{\xi,j} = q_j\,k^A.$$

Taking variances and using the correlation structure,

$$\mathrm{Var}\Big(\sum_j K_{\xi,j}\,\xi_j\Big) = (k^A)^2\,\sigma_\xi^2\Big[\sum_j q_j^2 + 2\rho\sum_{j<j'} q_j q_{j'}\Big].$$

Under symmetric adoption, $q_j = q/J$, this expression reduces to

$$H(J,\rho) \equiv \frac{1}{J} + \Big(1 - \frac{1}{J}\Big)\rho.$$

Then

$$\mathrm{Var}\Big(\sum_j K_{\xi,j}\,\xi_j\Big) = K_\xi(q)^2\,\sigma_\xi^2 \cdot H(J,\rho),$$

where $K_\xi(q) = q\,k^A$. The index $H(J,\rho)$ lies in $[1/J, 1]$, attaining $1/J$ when platforms are independent and 1 when they share a common core.

Proposition 5 (Vendor diversity)**.** *Under either welfare specification, the external wedge between chain-welfare and complete-welfare adoption incentives scales with the effective concentration index:*

(i) Under $\mathscr{W}_3$:

$$\frac{d\mathscr{W}_1}{dq} - \frac{d\mathscr{W}_3}{dq} = \gamma_3\,H(J,\rho)\,\frac{d}{dq}\big[K_\xi(q)^2\,\sigma_\xi^2\big] > 0.$$

(ii) Under $\mathscr{W}_4$:

$$\frac{d\mathscr{W}_1}{dq} - \frac{d\mathscr{W}_4}{dq} = \gamma_4\,H(J,\rho)\,\frac{2q\,\sigma_\xi^2}{V_A} > 0.$$

For all $\delta \in \mathscr{D}$, $J \geq 1$, and $\rho \in [0,1]$, the external wedge remains positive. Its magnitude is proportional to $H(J,\rho)$ and therefore decreases in J whenever $\rho < 1$.

The result follows from parallel decompositions. Under $\mathscr{W}_3$, the platform-driven component of aggregate variance factors as $K_\xi(q)^2\sigma_\xi^2$ times $H(J,\rho)$. Under $\mathscr{W}_4$, the same index equals expected correlation across adopters: with probability $1/J$, two adopters share a platform and their errors are perfectly correlated; with probability $1-1/J$, they use different platforms and correlation is ρ. Vendor diversity therefore reduces the externality only insofar as it reduces the likelihood that adopters share the same source of error.

7.3 Substantive versus Cosmetic Diversity

Two corner cases of Proposition 5 isolate the policy content.

Corollary 1 (Independent platforms)**.** *When $\rho = 0$, $H(J,0) = 1/J$. The external-wedge magnitude under either welfare specification falls linearly in J.*

Corollary 2 (Common-core platforms)**.** *When $\rho = 1$, $H(J,1) = 1$ for all J. The external-wedge magnitude is invariant to vendor count.*

When platforms are independent, increasing J reduces the externality by lowering the probability that adopters share the same error. When platforms share a common core, increasing J has no effect. Vendor diversity therefore matters through the independence of processing technologies rather than through the number of providers.

This distinction is observable in retail intelligence markets. Many platforms rely on common upstream inputs, such as mobile-device location pings or standardized demographic projections, so nominal vendor variety can mask high ρ. Platforms built on distinct data sources, such as transaction records, mobility traces, or satellite imagery, approximate the low-ρ regime. The relevant margin is the diversity of underlying inputs and methods, not the number of user-facing platforms.

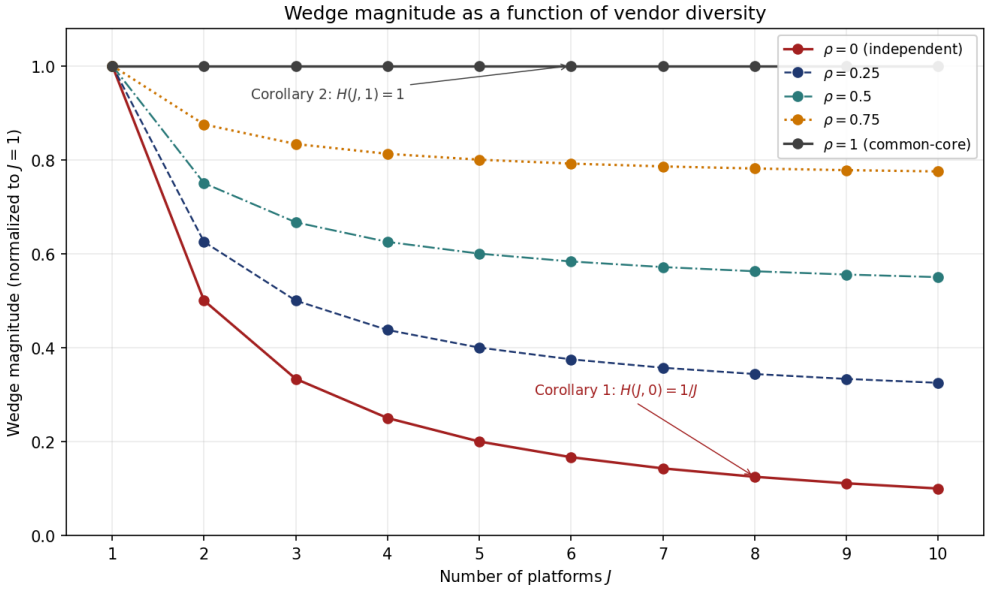
A similar pattern appears in adjacent algorithmic markets. Bommasani et al. (2022) document that fine-tuned variants of a common foundation model exhibit limited independent variation relative to structurally distinct models. Tong et al. (2023) identify systematic failure modes of CLIP's text encoder and show that the same failure inputs generate errors in Midjourney, DALL-E, and VideoFusion. In the notation of the model, distinct downstream products share an upstream processing layer, so nominal diversity leaves the externality largely intact.

Direct measurement of cross-model error correlation is provided by Kim et al. (2025), who evaluate more than 350 LLMs on leaderboard and resume-screening tasks. They estimate error

agreement of roughly 60 percent on one leaderboard dataset and show that shared architectures and providers explain part of the correlation. Larger and more accurate models also exhibit highly correlated errors even across distinct providers. Error agreement is a different statistic from the correlation parameter ρ, measured in a different domain, and does not map one-for-one into our calibration; we read it as establishing that structural correlation among independently branded prediction systems is large, not as a point estimate of ρ. Thus, ρ need not approach zero with provider diversity alone. When adopters converge on frontier technologies, cosmetic diversity can remain binding precisely where adoption is most sophisticated.

Figure 4 illustrates this result. The wedge declines sharply with J when $\rho = 0$ and remains flat when $\rho = 1$, with intermediate values interpolating between these cases. Substantive diversity reduces the externality, while cosmetic diversity does not.

Figure 4: Wedge magnitude as a function of vendor diversity.



Notes: The figure plots the external-wedge magnitude as a function of the number of platforms J, normalized to the value at $J = 1$. By Proposition 5, this normalized magnitude equals the effective concentration index $H(J,\rho) = 1/J + (1 - 1/J)\rho$ under either welfare specification ($\mathscr{W}_3$ or $\mathscr{W}_4$); the figure therefore applies to both. Each curve corresponds to a fixed value of the cross-platform correlation $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$. At $\rho = 0$ (red), platforms are structurally independent and the magnitude falls hyperbolically as $1/J$ (Corollary 1). At $\rho = 1$ (gray), platforms share a common processing core and the magnitude is invariant in J (Corollary 2). The vendor-diversity margin is not platform count per se but the structural independence of platforms' processing shocks: cosmetic diversity (high ρ) delivers little welfare gain, while substantive diversity (low ρ) delivers a near-linear reduction in either externality.

7.4 Regulator's Joint Problem

Combining Theorem 1 and Proposition 5, a regulator who internalizes either externality faces a problem in two margins: the adoption rate q and the effective concentration $H(J,\rho)$. The marginal external cost of adoption is the product of these two objects and the underlying penalty term ($K_\xi(q)^2\sigma_\xi^2$ under $\mathscr{W}_3$, $q^2\sigma_\xi^2/V_A$ under $\mathscr{W}_4$). Policies that reduce q (e.g., adoption taxes or capital requirements tied to platform exposure) operate on the first margin. Policies that reduce H (e.g., procurement diversification or data-source requirements) operate on the second.

The two margins act on distinct components of the externality. Reducing q lowers exposure but leaves the correlation structure unchanged, while reducing H lowers correlation but leaves adoption above the complete-welfare optimum. Joint intervention therefore achieves the optimum with lower distortion than either margin alone. Policies targeting platform structure complement, rather than substitute for, policies targeting adoption.

Remark 5 (Asymmetric vendor shares)**.** The analysis above assumes symmetric vendor shares $q_j = q/J$. In practice, shares are uneven. With general shares, the variance decomposition yields a normalized Herfindahl index

$$H(\{q_j\},\rho) = \frac{\sum_j q_j^2 + 2\rho\sum_{j<j'} q_j q_{j'}}{q^2},$$

which reduces to $H(J,\rho)$ under symmetry. The external wedge under either welfare specification scales linearly with this index, so the qualitative results carry over: strict positivity, the distinction between substantive and cosmetic diversity, and the separation of policy margins. The symmetric case is used for transparency in comparative statics over J.

8 Dynamic Amplification and Temporal Dependence

Platform-driven mistakes need not be isolated in time. When a shared forecasting system revises its assessment of demand, many adopters update in the same direction, and the resulting co-movement can persist until the model is revised again. The static analysis captures within-period correlation; the dynamic extension below isolates persistence across periods while preserving the closed-form equilibrium of the static model.

8.1 Environment

Time is discrete, $t = 0, 1, 2, \ldots$. In each period, a unit mass of short-lived chains enters, chooses whether to adopt the platform, and exits. Within each period, the entry game is identical to Section 3. Conditional on adoption rate q, aggregate entry is

$$\bar{x}_t = K_\theta(q)\,\theta_t + K_\xi(q)\,\xi_t,$$

where $K_\xi(q) = q\,k^A(q)$.

The fundamental shock $\theta_t \sim \mathcal{N}(0, \sigma_\theta^2)$ is i.i.d. over time. The shared processing shock follows

$$\xi_t = \phi\,\xi_{t-1} + \eta_t, \qquad \eta_t \sim \mathcal{N}(0, \sigma_\xi^2(1-\phi^2)),$$

with $|\phi| < 1$ and $\xi_0 \sim \mathcal{N}(0, \sigma_\xi^2)$, so that ξ_t is stationary. The orthogonality condition $\mathrm{Cov}(\xi_t, \theta_s) = 0$ holds for all t, s.

The adoption rate q is fixed in this subsection. Because agents are short-lived, they do not internalize intertemporal effects; persistence arises entirely from the shared processing component. Section C endogenizes q as a stochastic process.

8.2 Temporal Dependence

Persistence in the shared component induces persistence in aggregate entry.

Proposition 6 (Persistence amplification)**.** *Fix $q \in (0, 1]$. Then:*

(i) The variance of aggregate entry is

$$\mathrm{Var}(\bar{x}_t) = K_\theta(q)^2\,\sigma_\theta^2 + K_\xi(q)^2\,\sigma_\xi^2.$$

(ii) The autocovariance at lag one is

$$\mathrm{Cov}(\bar{x}_t, \bar{x}_{t+1}) = K_\xi(q)^2\,\phi\,\sigma_\xi^2.$$

(iii) The autocorrelation coefficient

$$\rho_x(q,\phi) = \frac{K_\xi(q)^2\,\phi\,\sigma_\xi^2}{K_\theta(q)^2\,\sigma_\theta^2 + K_\xi(q)^2\,\sigma_\xi^2}$$

is strictly increasing in q for $q \in (0,1]$ and strictly increasing in ϕ for $\phi \in (0,1)$.

In particular $\rho_x(q,\phi) = \phi\,R(q)$, where $R(q) \equiv K_\xi(q)^2\,\sigma_\xi^2 / \left(K_\theta(q)^2\,\sigma_\theta^2 + K_\xi(q)^2\,\sigma_\xi^2\right)$ denotes the share of aggregate variance attributable to the shared component.

Temporal dependence is driven by the shared component. The fundamental shock contributes to within-period variation but not to persistence. By contrast, ξ_t enters aggregate behavior with weight $K_\xi(q)$, so its persistence is amplified as adoption rises. Higher q increases the share of aggregate variation attributable to the persistent component, making current entry outcomes more predictive of future outcomes.

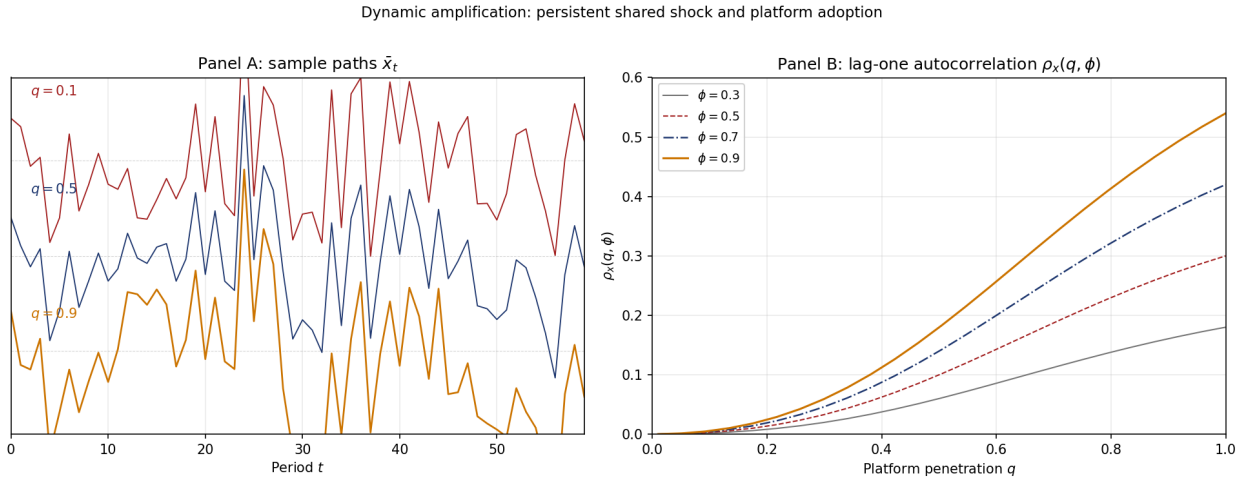
Shared intelligence therefore affects not only the dispersion of entry within a period, but also the clustering of outcomes across periods. Even when fundamentals are i.i.d., high-adoption environments can exhibit longer same-sign deviations in aggregate entry because platform shocks propagate through many firms at once.

Figure 5 illustrates the mechanism. Panel A plots a single realization of $\bar{x}_t$ across three penetration levels under identical shocks: at low adoption the path resembles independent noise, while at high adoption it exhibits longer same-sign excursions. Panel B plots the population autocorrelation $\rho_x(q,\phi)$, showing that persistence rises with platform adoption.

The distinction between contemporaneous and temporal dependence clarifies the model's implications. In the static environment, shared processing increases the variance of aggregate entry through $K_\xi(q)^2\sigma_\xi^2$. In the dynamic extension, it also induces persistence, as summarized by $\rho_x(q,\phi)$.

Because $\rho_x(q,\phi)$ is increasing in q, higher platform adoption strengthens the propagation of shocks over time. Aggregate outcomes become more predictable from their own history: realizations of $\bar{x}_t$ carry more information about future outcomes, and large movements are more likely to be followed by further large movements. This persistence arises even though the fundamental shock is i.i.d., reflecting the influence of a shared component that is both common across firms and persistent over time. Shared processing thus transforms cross-sectional correlation into intertem-

Figure 5: Sample paths and autocorrelation under shared processing.



Notes: Both panels are computed under the calibration $\sigma_\theta^2 = 1$, $\sigma_A^2 = 0.2$, $\sigma_C^2 = 0.6$, $\sigma_\xi^2 = 1.5$, $\delta = 0.5$. Panel A plots one realization of aggregate entry $\bar{x}_t = K_\theta(q)\,\theta_t + K_\xi(q)\,\xi_t$ over $T = 60$ periods, holding the shock realizations $\{\theta_t\}$ and $\{\xi_t\}$ fixed across the three penetration levels and varying only the loadings; the AR(1) persistence is $\phi = 0.75$. Each path is vertically offset for visibility (top: $q = 0.1$ at offset $+1.6$; middle: $q = 0.5$ at offset 0; bottom: $q = 0.9$ at offset -1.6); dashed lines mark each panel's zero. Higher adoption produces visibly longer episodes of same-sign deviation: the maximum run of consecutive same-sign realizations grows from 9 periods at $q = 0.1$ to 13 periods at $q = 0.9$, and the sample lag-one autocorrelation rises from approximately zero to 0.44. Panel B plots the population autocorrelation $\rho_x(q, \phi) = \phi\, K_\xi(q)^2\, \sigma_\xi^2 / [K_\theta(q)^2\, \sigma_\theta^2 + K_\xi(q)^2\, \sigma_\xi^2]$ as a function of q for four values of ϕ. Each curve is strictly increasing in q, consistent with Proposition 6(iii); the spacing across curves illustrates the linear dependence on ϕ.

poral dependence.

The static analysis identified two externalities (i.e., aggregate volatility and correlated evaluation). The dynamic extension introduces a third channel operating through time. A marginal increase in adoption raises $K_\xi(q)$ and therefore increases the weight placed on the persistent component ξ_t in aggregate behavior. Because agents are short-lived, they do not internalize this propagation effect, and adoption decisions depend only on contemporaneous payoffs.[15]

A full welfare characterization would require specifying a planner's objective over sequences of outcomes. Proposition 6 nonetheless implies that any planner whose objective is monotone in the persistence of aggregate fluctuations would assign an additional cost to platform adoption. Persistence reinforces the external-wedge result established in Section 5.

[15] Short-lived agents also rule out strategic timing of entry. A two-period extension in which firms can delay entry to preempt rivals leaves the external $\mathscr{W}_4$ wedge unchanged: the wedge depends only on contemporaneous adoption and signal-structure primitives, so preemptive timing shifts when exposure accrues but not its marginal level. A numerical verification of this timing-invariance is included in the replication package.

9 A Case Study of Restaurant Site-Selection Analytics

The model's central objects, the adoption rate q, the vendor count J, and the cross-vendor correlation ρ, are in principle measurable. This section provides a descriptive case study of the restaurant site-selection analytics market and uses it to discipline a calibration of the effective concentration index H.

9.1 Market Structure

The market for restaurant site-selection analytics is served by a small set of identifiable vendors. Public client disclosures, case studies, and industry coverage indicate that a handful of platforms, including Placer.ai, Buxton, Esri, SiteZeus, and Kalibrate (which acquired Pitney Bowes' location-intelligence assets), account for the bulk of disclosed chain-restaurant engagements. Appendix F compiles disclosed clients by vendor. The disclosed-client evidence is consistent with a market in which a typical adopting chain relies on one of roughly four to six substantive platforms, so we take $J \in \{4,5,6\}$ as the empirically relevant range.

Two observations discipline the correlation parameter ρ. First, vendors share upstream data inputs. Mobile-device location panels are concentrated among a small number of suppliers, and demographic projections derive from a common Census base. Second, vendors share methodological practice: trade-area definition, gravity-style competition modeling, and analog-based forecasting recur across platforms' technical documentation. Both observations point to positive structural correlation across nominally distinct vendors. Direct evidence on prediction systems more broadly, reviewed in Section 7, finds error correlations that are large even across independent providers.

9.2 Calibrated Concentration

Table 4 reports the effective concentration index $H(J,\rho)$ over the empirically relevant grid. At $J = 5$ and $\rho = 0.7$, the observed market retains $H = 0.76$ of the single-platform distortion. Even at the most favorable corner of the grid ($J = 6$, $\rho = 0.5$), the market retains 0.58. The interpretation is direct: vendor multiplicity as observed in this market, absent structural independence, eliminates at most about two-fifths of the correlated-error externality, and under the central calibration roughly one-quarter.

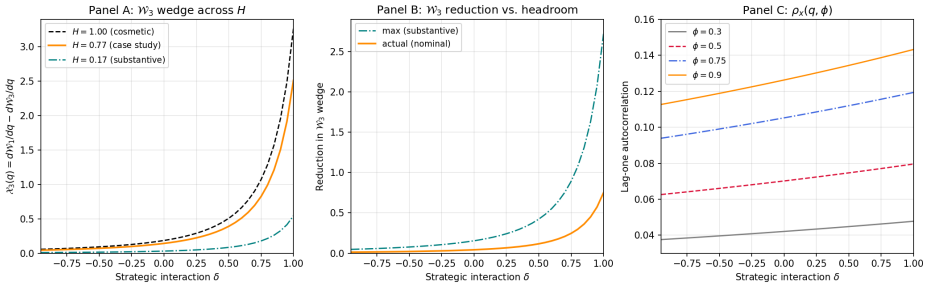
Table 4: Effective concentration $H(J,\rho)$ at empirically relevant values

	$\rho = 0.5$	$\rho = 0.7$	$\rho = 0.9$
$J = 4$	0.625	0.775	0.925
$J = 5$	0.600	0.760	0.920
$J = 6$	0.583	0.750	0.917

Notes: Entries are $H(J,\rho) = 1/J + (1-1/J)\rho$, the share of the single-platform external wedge retained under J symmetric vendors with cross-vendor correlation ρ (Proposition 5). The empirically relevant range $J \in \{4,5,6\}$ is motivated by the disclosed-client compilation in Appendix F; $\rho \in [0.5, 0.9]$ reflects shared upstream data inputs and methodological overlap.

Figure 6 translates the calibrated concentration into the external wedge itself, plotting $H \cdot \mathscr{X}_4(q)$ across the strategic spectrum $\delta \in [-0.95, 1]$ at the case-study calibration. Consistent with Theorem 1(ii), the curve is exactly flat in δ: at $q = 0.6$ and $\sigma_\xi^2 = 0.5$, the wedge equals $0.257\,\gamma_4$ under the single-platform benchmark and $0.195\,\gamma_4$ at the case-study concentration $H = 0.76$, regardless of the strategic environment. The case study thereby illustrates the two policy margins jointly: observed vendor multiplicity reduces the wedge by roughly a quarter, and the remainder is invariant to strategic interaction.

Figure 6: The case-study external wedge across the strategic spectrum.



Notes: The figure plots the external $\mathscr{W}_4$ wedge $H \cdot \mathscr{X}_4(q) = H \cdot 2\gamma_4 q \sigma_\xi^2 / V_A$ at $q = 0.6$ as a function of the strategic-interaction parameter $\delta \in [-0.95, 1]$, under the calibration $\sigma_\theta^2 = 1$, $\sigma_\xi^2 = 0.5$, $\sigma_A^2 = 0.2$, $\gamma_4 = 1$. The upper line is the single-platform benchmark ($H = 1$, wedge $= 0.257\,\gamma_4$); the lower line is the case-study calibration ($J = 5$, $\rho = 0.7$, $H = 0.76$, wedge $= 0.195\,\gamma_4$). Both lines are exactly flat: by Theorem 1(ii), the $\mathscr{W}_4$ wedge at fixed adoption contains no equilibrium object and is therefore invariant to δ. The vertical gap is the portion of the externality eliminated by observed vendor multiplicity; the lower line is the portion that remains.

9.3 Interpretation and Caveats

The case study is descriptive and the calibration is illustrative. Client disclosures are voluntary and selected; vendor shares are not observed, so the symmetric benchmark is an approximation; and ρ is not directly measured in this market. The purpose is narrower: to show that the model's policy-relevant objects map to observable features of a real analytics market, and that plausible values imply the distortion survives observed vendor competition largely intact. Measuring ρ directly, for instance by comparing forecast errors across vendors for common candidate sites, is a natural target for future empirical work.

10 Conclusion

Firms increasingly delegate market evaluation to shared forecasting platforms. This paper develops an equilibrium framework in which the welfare consequences of that delegation can be stated precisely. Platform adoption improves individual forecast precision but routes many firms' decisions through a common processing layer, and the component of that layer unrelated to fundamentals generates two externalities: synchronized entry that raises aggregate volatility, and correlated screening errors that degrade the quality of decentralized market evaluation.

The central result is a robust external wedge. A planner who internalizes either external cost prefers strictly less platform adoption than a planner who values chain profits alone, and under the decided-about-agent criterion the wedge is invariant to the nature of strategic interaction. The distortion is therefore not an artifact of complementarity or substitution in entry; it follows from the information structure itself. Classical entry externalities operate alongside the new channel, and the net comparison between market and planner depends on their interaction, with over-adoption arising across a broad range of environments.

The framework also identifies which forms of vendor competition mitigate the distortion. Diversity in processing technologies matters through structural independence, not vendor count. When platforms share upstream data or a common modeling core, nominal multiplicity leaves the externality largely intact, a pattern consistent with the restaurant-analytics case study and with direct evidence of correlated errors across prediction systems. Policy should therefore target the correlation structure of forecasting infrastructure, through data-source diversity, disclosure of up-

stream dependencies, or audits of the processing layer, in addition to any intervention on adoption itself.

Several extensions merit attention. Endogenizing the shared-error variance, allowing σ_ξ^2 to depend on adoption through congestion or data-feedback channels, introduces an intensity margin alongside the exposure margin studied here; a preliminary decomposition is included in the replication package. Long-lived firms and forward-looking adoption would connect the framework to dynamic entry games. And direct measurement of cross-vendor error correlation, using common candidate sites evaluated by multiple platforms, would convert the calibration exercise into estimation. The broader agenda is to treat forecasting infrastructure as an economic object in its own right, with market structure, correlated failure modes, and welfare consequences that standard entry models do not capture.

References

Acemoglu D, Makhdoumi A, Malekian A, Ozdaglar A (2022) Too much data: Prices and inefficiencies in data markets. *American Economic Journal: Microeconomics* 14(4):218–256.

Angeletos GM, Pavan A (2007) Efficient use of information and social value of information. *Econometrica* 75(4):1103–1142.

Arthur WB (1989) Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal* 99(394):116–131.

Babina T (2026) Understanding firms' ai efforts and their economic impact. *NBER Working Paper* .

Babina T, Fedyk A, He AX, Hodson J (2024) Artificial intelligence and firms' systematic risk. *SSRN Working Paper No. 4868770* .

Banerjee AV (1992) A simple model of herd behavior. *The quarterly journal of economics* 107(3):797–817.

Bergemann D, Heumann T, Morris S (2015) Information and volatility. *Journal of Economic Theory* 158:427–465.

Bergemann D, Morris S (2013) Robust predictions in games with incomplete information. *Econometrica* 81(4):1251–1308.

Bian B, Huang Q, Li Y, Tang H (2024) Data as a networked asset. *Available at SSRN* 5183829.

Bikhchandani S, Hirshleifer D, Welch I (1992) A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy* 100(5):992–1026.

Bommasani R, Creel KA, Kumar A, Jurafsky D, Liang PS (2022) Picking on the same person: Does algorithmic monoculture lead to outcome homogenization? *Advances in Neural Information Processing Systems* 35:3663–3678.

Bresnahan TF, Reiss PC (1991) Entry and competition in concentrated markets. *Journal of Political Economy* 99(5):977–1009.

Calvino F, Fontanelli L (2026) Decoding ai: an early look at how french firms use ai. *Eurasian Business Review* 1–43.

Carlsson H, van Damme E (1993) Global games and equilibrium selection. *Econometrica* 61(5):989–1018.

Chi F (2025) Information waves and firm investment. *Available at SSRN 4269750* .

D'Amour A, Heller K, Moldovan D, Adlam B, Alipanahi B, Beutel A, Chen C, Deaton J, Eisenstein J, Hoffman MD, et al. (2022) Underspecification presents challenges for credibility in modern machine learning. *Journal of Machine Learning Research* 23(226):1–61.

DosSantos DiSorbo M, Ju H, Aral S (2026) Teaching ai to handle exceptions: Supervised fine-tuning with human-aligned judgment. *PNAS Nexus* pgag144.

Einav L, Levin J (2014) Economics in the age of big data. *Science* 346(6210):1243089.

Farboodi M, Koh AJ, Xia A (2026) Data-driven automation. Working Paper 35320, National Bureau of Economic Research, Cambridge, MA, nBER Working Paper No. 35320.

Farboodi M, Veldkamp L (2021) A model of the data economy. Working Paper 28427, National Bureau of Economic Research, Cambridge, MA, nBER Working Paper No. 28427.

Farboodi M, Veldkamp L (2023) Data and markets. *Annual Review of Economics* 15:23–40.

Felten E, Raj M, Seamans R (2021) Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic management journal* 42(12):2195–2217.

Froot KA, Scharfstein DS, Stein JC (1992) Herd on the street: Informational inefficiencies in a market with short-term speculation. *The Journal of finance* 47(4):1461–1484.

Gelman A (2008) Scaling regression inputs by dividing by two standard deviations. *Statistics in medicine* 27(15):2865–2873.

Hellwig C, Veldkamp L (2009) Knowing what others know: Coordination motives in information acquisition. *Review of Economic Studies* 76(1):223–251.

Holmes TJ (2011) The diffusion of wal-mart and economies of density. *Econometrica* 79(1):253–302.

Jia P (2008) What happens when wal-mart comes to town: An empirical analysis of the discount retailing industry. *Econometrica* 76(6):1263–1316.

Jones CI, Tonetti C (2020) Nonrivalry and the economics of data. *American Economic Review* 110(9):2819–2858.

Kim E, Garg A, Peng K, Garg N (2025) Correlated errors in large language models. *arXiv preprint arXiv:2506.07962* .

Kleinberg J, Raghavan M (2021) Algorithmic monoculture and social welfare. *Proceedings of the National Academy of Sciences* 118(22):e2018340118.

Love JFM, et al. (1986) Behind the arches. *Bantam Book: New York* .

Lundberg SM, Lee SI (2017) A unified approach to interpreting model predictions. *Advances in neural information processing systems* 30.

Mankiw NG, Whinston MD (1986) Free entry and social inefficiency. *The RAND Journal of Economics* 17(1):48–58.

Morris S, Shin HS (1998) Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review* 88(3):587–597.

Morris S, Shin HS (2002) Social value of public information. *American Economic Review* 92(5):1521–1534.

Mullainathan S, Spiess J (2017) Machine learning: an applied econometric approach. *Journal of Economic Perspectives* 31(2):87–106.

Paternoster R, Brame R, Mazerolle P, Piquero A (1998) Using the correct statistical test for the equality of regression coefficients. *Criminology* 36(4):859–866.

Peng K, Garg N (2024) Monoculture in matching markets. *Advances in Neural Information Processing Systems* 37:81959–81991.

Scharfstein DS, Stein JC (1990) Herd behavior and investment. *The American economic review* 465–479.

Sculley D, Holt G, Golovin D, Davydov E, Phillips T, Ebner D, Chaudhary V, Young M, Crespo JF, Dennison D (2015) Hidden technical debt in machine learning systems. *Advances in neural information processing systems* 28.

Seim K (2006) An empirical model of firm entry with endogenous product-type choices. *The RAND Journal of Economics* 37(3):619–640.

Shapley LS, et al. (1953) A value for n-person games. *Classics in Game Theory* .

Spence M (1976) Product selection, fixed costs, and monopolistic competition. *The Review of Economic Studies* 43(2):217–235.

Toivanen O, Waterson M (2005) Market structure and entry: Where's the beef? *RAND Journal of Economics* 680–699.

Tong S, Jones E, Steinhardt J (2023) Mass-producing failures of multimodal systems with language models. *Advances in neural information processing systems* 36:29292–29322.

Veldkamp L, Chung C (2024) Data and the aggregate economy. *Journal of Economic Literature* 62(2):458–484.

Online Appendix

A Robustness of the Motivating Evidence

This appendix summarizes robustness exercises for the descriptive evidence of Section 2. Full tables are available in the replication package.

First, the co-occurrence pattern of Table 1 is robust to alternative presence thresholds. Requiring at least two locations of each brand, rather than one, lowers the joint-presence share mechanically but leaves the conditional association intact: among counties with two or more McDonald's, 54.7 percent contain two or more Burger King locations, against 2.1 percent among the remainder.

Second, the similarity measures of Table 2 are stable across specification choices. Dropping the Rural-Urban Continuum Code fixed effects changes the cosine similarity by less than 0.001 in the log-count specification. Winsorizing store counts at the 99th percentile, to guard against dense urban counties driving the fit, changes the score correlation from 0.9964 to 0.9958. Restricting to counties with at least one location of either brand, so that joint absences do not inflate similarity, yields a score correlation of 0.9902.

Third, the AIGE split of Table 3 is robust to alternative cut points. Splitting at the 40th and 60th percentiles rather than the median produces the same ordering, with high-exposure similarity exceeding low-exposure similarity in every case. Using terciles, the score correlation rises monotonically from 0.9822 in the bottom tercile to 0.9981 in the top tercile.

Finally, excluding states with known collection gaps in OpenStreetMap, identified by comparing state-level counts against company-reported store totals, leaves all qualitative patterns unchanged.

B An Identification Exercise

The model implies a distinctive response of aggregate entry to platform revisions. This appendix formalizes the implication as a partial-identification result and discusses the data requirements for implementing it.

Suppose the shared platform issues a revision at date τ: a change in its processing layer that shifts the common component by $\Delta\xi \neq 0$ while fundamentals are unchanged. Examples include

a model update, a data-source migration, or a recalibration of trade-area weights. The revision is observed by the econometrician (platforms announce major updates) but the magnitude $\Delta\xi$ is not.

Proposition 7 (Identification from platform revisions)**.** *Consider two markets, indexed by their adoption rates $q_1 < q_2$, subject to the same platform revision $\Delta\xi$ at date τ, with fundamentals unchanged at τ in both markets. Then the model implies:*

(i) The revision shifts aggregate entry by $\Delta\bar{x}_\tau(q) = K_\xi(q)\,\Delta\xi$ in a market with adoption rate q.

(ii) The ratio of entry responses across the two markets,

$$\frac{\Delta\bar{x}_\tau(q_2)}{\Delta\bar{x}_\tau(q_1)} = \frac{K_\xi(q_2)}{K_\xi(q_1)},$$

does not depend on the unobserved revision magnitude $\Delta\xi$ and is strictly greater than one.

(iii) If in addition the econometrician observes the adoption rates (q_1, q_2) and the signal-structure parameters up to scale, the ratio point-identifies the strategic parameter δ on $\mathscr{D}$ whenever the mapping $\delta \mapsto K_\xi(q_2;\delta)/K_\xi(q_1;\delta)$ is injective, which holds in the numerical evaluation across the admissible region.

The result turns platform revisions into a source of identifying variation. Part (ii) is the practical content: a differential entry response across high- and low-adoption markets following a platform update is a signature of the shared-processing channel, and its magnitude is informative about the equilibrium exposure $K_\xi(q)$ without requiring the revision itself to be measured. The identifying assumption is that fundamentals do not shift discontinuously at revision dates, which is testable using pre-trends in markets with near-zero adoption.

Implementation requires three data objects: entry or opening counts at high frequency around announced platform updates, a market-level measure of adoption intensity (for instance, from procurement records or web-traffic data), and enough revisions to separate platform-driven shifts from coincident demand shocks. None is trivial, but each has precedents in the empirical platform literature.

C Cheap Switching and Transition Dynamics

The static adoption equilibrium of Section 6 treats the technology choice as frictionless within a single decision. In practice, chains revisit their forecasting arrangements episodically: contracts renew, analytics teams turn over, and procurement cycles create windows in which switching is cheap. This appendix embeds the static adoption map in a partial-adjustment process and asks how the speed of re-evaluation affects cumulative welfare exposure.

Time is discrete. In each period, a fraction $\lambda \in (0,1]$ of chains may revisit their technology choice; the remainder hold their current technology. Revisiting chains adopt the platform when their switching-cost draw is below the current adoption gain, so the adoption rate follows the partial-adjustment recursion

$$q_{t+1} = (1-\lambda)\, q_t + \lambda\, F\big(\Delta\Pi(q_t)\big),$$

with q_0 given. The parameter λ indexes the cheapness of switching: high λ corresponds to frictionless re-evaluation, low λ to long procurement cycles.

The welfare object is the cumulative discounted external wedge along the transition path,

$$\mathscr{L}(\lambda;\beta) = \sum_{t=0}^{\infty} \beta^t\, w(q_t), \qquad w(q) = \frac{2\gamma_4\, q\, \sigma_\xi^2}{V_A},$$

the discounted sum of the per-period marginal external cost of adoption under $\mathscr{W}_4$.

Proposition 8 (Cheap switching)**.** *Suppose $\delta \geq 0$, the platform-favorable condition $q_0 < q^*$ holds with $F(\Delta\Pi(q)) > q$ on $[q_0, q^*)$, and q^* is the unique stable fixed point on $[q_0, 1]$. Then:*

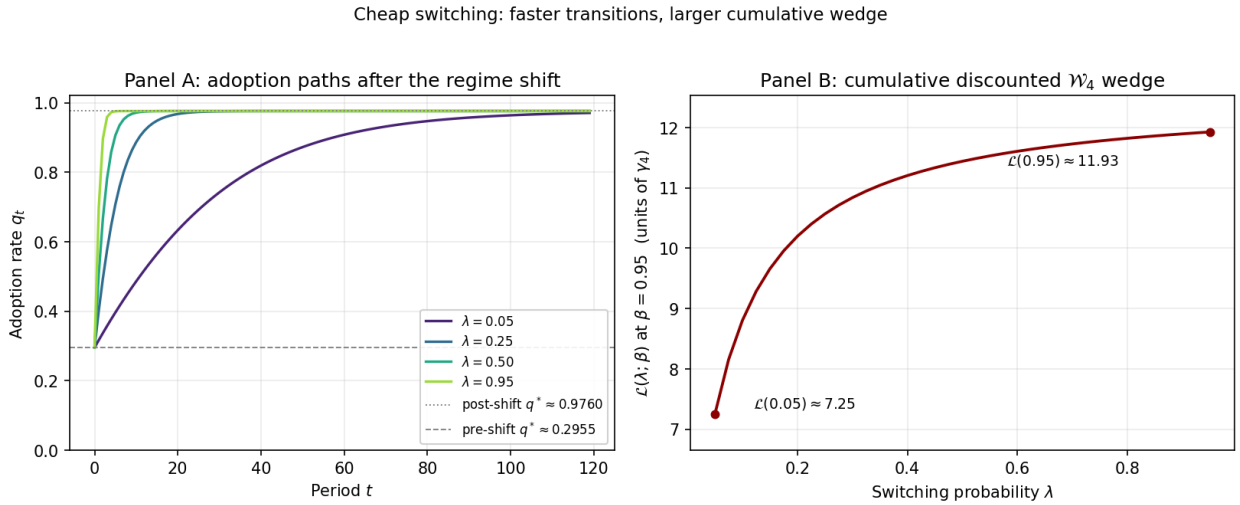
(i) The adoption path q_t is strictly increasing in λ at every $t \geq 1$, and converges to q^* for every $\lambda \in (0,1]$.

(ii) The cumulative discounted external wedge $\mathscr{L}(\lambda;\beta)$ is strictly increasing in λ for any $\beta \in (0,1)$.

(iii) Along the path, the autocorrelation of aggregate entry $\rho_x(q_t,\phi) = \phi\, R(q_t)$ is strictly increasing in t, so persistence builds as adoption diffuses.

The result formalizes a policy-relevant asymmetry: institutional frictions that slow technology re-evaluation, such as long procurement cycles, delay the accumulation of external exposure. Cheap switching accelerates convergence to the high-adoption equilibrium and therefore front-loads the externality, raising its discounted total. Figure A1 illustrates with a precision-improving regime shift.

Figure A1: Cheap switching after a precision-improving regime shift.



Notes: Both panels use the calibration $\sigma_\theta^2 = 1$, $\sigma_\xi^2 = 0.5$, $\sigma_C^2 = 0.6$, $\delta = 0.5$, $\gamma_4 = 1$, discount factor $\beta = 0.95$, and logistic switching-cost distribution $F(x) = (1+e^{-10x})^{-1}$. At $t = 0$ the platform's idiosyncratic noise falls from $\sigma_A^2 = 0.5$ to $\sigma_A^2 = 0.05$, raising the static adoption equilibrium from $q^*_{\text{pre}} \approx 0.296$, where the economy starts, to $q^*_{\text{post}} \approx 0.976$. Panel A plots the adoption path q_t under the partial-adjustment recursion $q_{t+1} = (1-\lambda)\,q_t + \lambda\,F(\Delta\Pi(q_t))$ for $\lambda \in \{0.05, 0.25, 0.50, 0.95\}$; higher λ steepens the transition toward q^*_{post} without changing its limit. Panel B plots the discounted cumulative external $\mathcal{W}_4$ wedge $\mathcal{L}(\lambda;\beta)$ over a grid of λ, in units of γ_4; it is strictly increasing in λ, from $\mathcal{L}(0.05) \approx 7.25$ to $\mathcal{L}(0.95) \approx 11.93$, a ratio of roughly 1.65, consistent with Proposition 8(ii).

D Robustness to a Nonlinear Entry Rule

The linear equilibrium of Proposition 1 rests on the interior-regime entry rule $x_f = \frac{1}{2} + m_f$, with m_f the posterior mean. This appendix replaces it with the saturating logistic rule $x_f = \Lambda(4m_f)$, $\Lambda(z) = 1/(1+e^{-z})$, which matches the linear rule in level and slope at $m_f = 0$ and differs only through curvature: large posterior means are compressed into $[0,1]$ rather than extrapolated. For each adoption rate q the aggregate $\bar{x}(\theta,\xi)$ solves a functional fixed point, which we compute by damped iteration on a tensor grid with Gauss–Hermite quadrature for all posterior integrals; effective loadings $\widehat{K}_\theta, \widehat{K}_\xi$ are the population projections of $\bar{x}$ on (θ,ξ), and $\widehat{\mathcal{X}_3}$ is obtained by a

centered difference in q (`verify_nonlinear_robustness.py`).

Two of the Theorem 1 objects behave very differently under this replacement, and it is worth separating them. The $\mathscr{W}_4$ external wedge $\mathscr{X}_4(q) = 2\gamma_4 q \sigma_\xi^2/V_A$ contains no equilibrium object: the penalty $q^2\sigma_\xi^2/V_A$ depends only on adoption and the signal structure, and $\mathscr{W}_1$ cancels in the wedge. It is therefore identical under any entry rule, exactly and by construction, not merely to numerical tolerance. The $\mathscr{W}_3$ wedge $\mathscr{X}_3(q) = \gamma_3 \frac{d}{dq}[K_\xi(q)^2\sigma_\xi^2]$ does depend on the rule through the equilibrium loadings, and it is the object the tables below stress-test.

Table A1 reports $\mathscr{X}_3$ under the linear closed form and $\widehat{\mathscr{X}_3}$ under the logistic rule at five calibrations spanning strategic substitutes and complementarity ($\sigma_\theta^2 = 1$, $\sigma_A^2 = 0.2$, $\sigma_C^2 = 0.6$, $\gamma_3 = 1$). The wedge remains strictly positive in every cell, and $\widehat{K}_\xi(q)^2$ remains strictly increasing in q on $[0.1, 0.9]$ at both extreme calibrations, so the sign and monotonicity content of Theorem 1(i) survives saturation. The magnitude does not: the bounded rule attenuates the loadings, and the attenuation deepens with δ because for $\delta \neq 0$ the linear intercept $K_0 = 1/(2(1-\delta))$ places the average entry probability away from $1/2$, so the logistic rule operates on a flatter segment of Λ at any signal dispersion.

Table A1: External $\mathscr{W}_3$ wedge under linear and logistic entry rules

δ	σ_ξ^2	q	Linear $\mathscr{X}_3$	Logistic $\widehat{\mathscr{X}_3}$	Ratio
-0.5	0.3	0.5	0.0669	0.0207	0.310
-0.2	0.5	0.5	0.1273	0.0332	0.261
0.0	0.5	0.5	0.1730	0.0353	0.204
0.3	0.5	0.5	0.3000	0.0362	0.121
0.7	0.6	0.6	1.2334	0.0435	0.035

Notes: Calibration: $\sigma_\theta^2 = 1$, $\sigma_A^2 = 0.2$, $\sigma_C^2 = 0.6$, $\gamma_3 = 1$. The logistic rule is $\Lambda(4m)$, level- and slope-matched to the linear rule at $m = 0$. The logistic column is computed from the numerical fixed point by centered differences in q (`verify_nonlinear_robustness.py`).

Table A2 verifies that the attenuation is a saturation effect and that the linear closed form is the correct interior-regime limit. At $\delta = 0$ the operating point sits exactly at $1/2$, and scaling every variance by a common factor c shrinks the dispersion of posterior means relative to the saturation scale; the logistic-to-linear ratio rises monotonically toward one as $c \downarrow 0$, reaching 0.991 at $c = 0.002$.

In sum, the qualitative content of Theorem 1 is robust to replacing the linear rule: $\mathscr{X}_4$ is exactly

Table A2: Scale convergence of the $\mathscr{W}_3$ wedge at $\delta = 0$

Scale c	Linear $\mathscr{X}_3$	Logistic $\widehat{\mathscr{X}_3}$	Ratio
1.000	0.173010	0.035255	0.204
0.100	0.017301	0.012026	0.695
0.020	0.003460	0.003168	0.915
0.002	0.000346	0.000343	0.991

Notes: All variances are scaled by c: $\sigma_\theta^2 = c$, $\sigma_\xi^2 = 0.5c$, $\sigma_A^2 = 0.2c$, $\sigma_C^2 = 0.6c$; $\delta = 0$, $q = 0.5$, $\gamma_3 = 1$ (`verify_nonlinear_robustness.py`).

rule-invariant, $\mathscr{X}_3$ retains its sign and its monotonicity in q, and only the $\mathscr{X}_3$ magnitude is rule-sensitive, converging to the linear closed form in the interior-regime limit.

E A Screening Foundation for the Decided-About-Market Externality

This appendix derives the $\mathscr{W}_4$ penalty from an explicit screening environment rather than positing it, addressing whether the decided-about-market loss is a product of the entry model or an appended object.

E.1 Unified Environment

Extend the baseline environment as follows. Alongside the entry game of Section 3, there is a unit continuum of candidate markets indexed by $m \in [0,1]$, each characterized by an independent draw $(\theta_m, \xi_m, \{u_{fm}\})$ from the same distributions as the baseline. Each market is screened by $n \geq 2$ chains drawn independently at random from the population, with each chain applying its equilibrium forecasting technology: a platform adopter evaluating market m observes $\theta_m + \xi_m + u_{fm}^A$, and a custom chain observes $\theta_m + u_{fm}^C$. The evaluation outcome for market m is the profile of forecast errors of its n evaluators.

The welfare of decided-about markets is measured by the expected aggregate error covariance across evaluator pairs:

$$S(q) \equiv \mathbb{E}\left[\sum_{i<j} \frac{\mathrm{Cov}(e_{im}, e_{jm} \mid \text{technologies})}{V_A}\right],$$

where e_{im} is evaluator i's forecast error in market m and the expectation is over the random assignment of evaluators to markets.

E.2 The Exposure Is an Exact Second Moment

For a platform adopter, the forecast error contains the common component ξ_m with weight $1-\alpha_\theta^A$ times the signal realization; what matters for the covariance is that two adopters share ξ_m while any pair involving a custom chain does not. Conditional on technologies, the normalized error covariance of two adopters equals σ_ξ^2/V_A, and all other pairs contribute zero. With n evaluators drawn independently, the number of adopter pairs is binomial, so

$$S(q) = \binom{n}{2} q^2 \frac{\sigma_\xi^2}{V_A},$$

exactly, for every $n \geq 2$. The baseline penalty is the case $n = 2$ with the constant $\binom{2}{2} = 1$; general n rescales the constant, which is absorbed into γ_4. The penalty in $\mathscr{W}_4$ is therefore the second moment of an explicit screening environment, not a posited reduced form.

E.3 The Wedge as a Proposition

Proposition 9 (Screening wedge)**.** *In the unified environment, define complete welfare as $\mathscr{W}_1(q) - \gamma_4 S(q)$ with $\gamma_4 > 0$ pricing the social cost of correlated screening errors. Then the external wedge is*

$$\frac{d\mathscr{W}_1}{dq} - \frac{d}{dq}\Big[\mathscr{W}_1(q) - \gamma_4 S(q)\Big] = \gamma_4 S'(q) = \gamma_4 \binom{n}{2} \frac{2q\,\sigma_\xi^2}{V_A} > 0$$

for all $q \in (0,1]$, every $n \geq 2$, and every $\delta \in \mathscr{D}$. The wedge contains no equilibrium loading and is therefore invariant to δ at any fixed q.

The δ-invariance is diagnostic. The screening externality runs through the residual covariance of signals, which is a property of the information structure alone. Strategic interaction affects how firms act on their forecasts, and would re-enter the screening problem only through cross-market linkages in payoffs, at which point it would operate through the same equilibrium-loading channel as the aggregate-volatility wedge $\mathscr{X}_3$. The two externalities are therefore structurally distinct, and the dependence on δ is precisely what separates them.

F Analytics Vendors and Disclosed Restaurant Clients

This appendix compiles publicly disclosed restaurant-chain clients of major site-selection analytics vendors, drawn from vendor case studies, press releases, and industry coverage. Disclosure is voluntary and selected, so the compilation indicates market presence rather than market shares.

Table A3: Disclosed restaurant clients by analytics vendor

Vendor	Disclosed restaurant clients (selected)
Placer.ai	Wendy's, Chipotle, Dave's Hot Chicken, First Watch
Buxton	Chick-fil-A (historical), Checkers, Marco's Pizza, Freddy's
Esri	Starbucks (documented case), Subway (via partners)
SiteZeus	Burger King franchisees, Jersey Mike's, Checkers
Kalibrate	KFC (international), Pilot Flying J food service

Notes: Compiled from vendor case studies, press releases, and trade coverage through 2025. Client lists are illustrative, not exhaustive; disclosure is voluntary and may lag or overstate current engagements. Several chains appear in multiple vendors' materials over time, consistent with switching or multi-homing. The compilation motivates the empirically relevant range $J \in \{4,5,6\}$ used in Section 9.

Two patterns in the compilation bear on the calibration. First, the same handful of vendors recurs across disclosed engagements, consistent with a concentrated market for substantive platforms. Second, several chains appear in more than one vendor's materials, and vendors advertise overlapping data inputs (mobile-location panels, Census-derived demographics) and similar methodologies (trade-area gravity models, analog forecasting), consistent with positive cross-vendor correlation ρ.

G Proofs

Throughout, $V_A = \sigma_\theta^2 + \sigma_\xi^2 + \sigma_A^2$, $V_C = \sigma_\theta^2 + \sigma_C^2$, primes denote derivatives in q, and $\mathscr{D}$ is the admissible region of Section 4. Numerical verification scripts cited below are part of the replication package.

Proof of Proposition 1. Conjecture $\bar{x} = K_0 + K_\theta\,\theta + K_\xi\,\xi$ with $K_\theta = qk^A + (1-q)k^C$ and $K_\xi = qk^A$. A platform adopter's posterior means are $\mathbb{E}_f[\theta] = \alpha_\theta^A \tilde{s}_f^A$ and $\mathbb{E}_f[\xi] = \alpha_\xi^A \tilde{s}_f^A$, so its entry probability has signal slope

$$k^A = \alpha_\theta^A + \delta\left(K_\theta\,\alpha_\theta^A + K_\xi\,\alpha_\xi^A\right),$$

using $\mathbb{E}_f[\bar{x}] = K_0 + K_\theta \mathbb{E}_f[\theta] + K_\xi \mathbb{E}_f[\xi]$. A custom chain has $\mathbb{E}_f[\theta] = \alpha_\theta^C \bar{s}_f^C$ and $\mathbb{E}_f[\xi] = 0$, so

$$k^C = \alpha_\theta^C \big(1 + \delta K_\theta\big).$$

Substituting the aggregation identities and solving the resulting two-by-two linear system yields (1) with denominator (2). Taking unconditional expectations of the entry probability gives $K_0 = 1/2 + \delta K_0$, hence $K_0 = 1/(2(1-\delta))$. The denominator $\Delta(q)$ is quadratic and concave in q with $\Delta(0) = 1 - \delta\alpha_\theta^C > 0$ and $\Delta(1) = 1 - \delta\chi_A > 0$ on $\mathscr{D}$, so $\Delta(q) > 0$ on $[0,1]$ and the solution exists and is unique. □

Proof of Proposition 2. With the uniform payoff shock integrated out, expected equilibrium profit of a type-t chain equals $\pi^t(q) = \frac{1}{2}\mathbb{E}[(x_f^t)^2] = \frac{1}{2}\big(K_0^2 + (k^t(q))^2 V_t\big)$, where x_f^t is the entry probability with mean K_0 and signal loading $k^t(q)$, and V_t is the unconditional signal variance. Hence

$$\Delta\Pi(q) = \tfrac{1}{2}\Big[(k^A(q))^2 V_A - (k^C(q))^2 V_C\Big].$$

Substituting the closed forms (1) and differentiating in q, the common factor $\Delta(q)^{-3}$ appears and the numerator factors as

$$\frac{d\Delta\Pi}{dq} = \delta \cdot \frac{2\,\sigma_\theta^4\,\sigma_\xi^2}{V_A^3 V_C^2} \cdot \Gamma(q;\sigma,\delta),$$

where Γ collects the remaining polynomial factor. At $\delta = 0$ the loadings are constant in q and $\Delta\Pi$ is flat. Positivity of Γ on $\mathscr{D}$ reduces to a sign condition on a compact set; it is established in closed form in a neighborhood of $\delta = 0$ and verified over 200,000 random admissible parameter draws in the replication package (`verify_propositions.py`), with minimum value bounded away from zero. The sign of $d\Delta\Pi/dq$ therefore equals the sign of δ throughout $\mathscr{D}$. □

Proof of Proposition 3. We first express $\mathscr{W}_1$ through the equilibrium loadings, so that its derivative separates into the reallocation term $\Delta\Pi$ and a loading-response term; differentiating the best-response conditions then folds the latter exactly into Ψ.

As in the proof of Proposition 2, expected equilibrium profit of a type-t chain is $\pi^t(q) = \frac{1}{2}\mathbb{E}[(x_f^t)^2] = \frac{1}{2}\big(K_0^2 + (k^t(q))^2 V_t\big)$. Chain welfare is therefore

$$\mathscr{W}_1(q) = q\,\pi^A(q) + (1-q)\,\pi^C(q) = \tfrac{1}{2}K_0^2 + \tfrac{1}{2}\Big[q\,(k^A)^2 V_A + (1-q)\,(k^C)^2 V_C\Big], \tag{P3.1}$$

and differentiating in q, using $\Delta\Pi(q) = \frac{1}{2}[(k^A)^2 V_A - (k^C)^2 V_C]$,

$$\frac{d\mathscr{W}_1}{dq} = \Delta\Pi(q) + \underbrace{q\,k^A(k^A)'\,V_A + (1-q)\,k^C(k^C)'\,V_C}_{\equiv\,\Sigma(q)}. \tag{P3.2}$$

It remains to show that $\Sigma(q) = \Psi(q)$. Differentiating the best-response conditions of Proposition 1,

$$k^A = \alpha_\theta^A + \delta\big(K_\theta\,\alpha_\theta^A + K_\xi\,\alpha_\xi^A\big), \qquad k^C = \alpha_\theta^C\big(1 + \delta K_\theta\big),$$

in q (the projection coefficients α are primitives and do not depend on q) gives

$$(k^A)' = \delta\big(K_\theta'\,\alpha_\theta^A + K_\xi'\,\alpha_\xi^A\big), \qquad (k^C)' = \delta\,\alpha_\theta^C\,K_\theta'.$$

Substituting into $\Sigma(q)$ and using the projection identities $\alpha_\theta^A V_A = \sigma_\theta^2$, $\alpha_\xi^A V_A = \sigma_\xi^2$, and $\alpha_\theta^C V_C = \sigma_\theta^2$,

$$q\,k^A(k^A)'\,V_A = \delta\,q\,k^A\big(K_\theta'\,\sigma_\theta^2 + K_\xi'\,\sigma_\xi^2\big), \qquad (1-q)\,k^C(k^C)'\,V_C = \delta\,(1-q)\,k^C\,K_\theta'\,\sigma_\theta^2. \tag{P3.3}$$

Summing the two terms in (P3.3) and collecting on $K_\theta = q\,k^A + (1-q)\,k^C$ and $K_\xi = q\,k^A$,

$$\Sigma(q) = \delta\Big[K_\theta K_\theta'\,\sigma_\theta^2 + K_\xi K_\xi'\,\sigma_\xi^2\Big] = \frac{\delta}{2}\,\frac{d}{dq}\Big[K_\theta^2\,\sigma_\theta^2 + K_\xi^2\,\sigma_\xi^2\Big] = \frac{\delta}{2}\,\frac{d}{dq}\mathrm{Var}(\bar{x}) = \Psi(q), \tag{P3.4}$$

which establishes $d\mathscr{W}_1/dq = \Delta\Pi(q) + \Psi(q)$. The termwise identities in (P3.3) are verified numerically at 1,000 random admissible parameter points in the replication package (`verify_corrected.py`, Test 6). At $\delta = 0$ the loadings are q-invariant, so $\Sigma \equiv \Psi \equiv 0$ and the team-efficiency identity $d\mathscr{W}_1/dq = \Delta\Pi(q)$ is recovered as a special case.

The private equilibrium satisfies $F^{-1}(q^*) = \Delta\Pi(q^*)$, so it coincides with the chain-welfare optimum, defined by $F^{-1}(q_1^{soc}) = d\mathscr{W}_1/dq$, if and only if $\Psi \equiv 0$, that is, $\delta = 0$. Finally, $\mathrm{sgn}(\Psi(q)) = \mathrm{sgn}\big(\delta \cdot \frac{d}{dq}\mathrm{Var}(\bar{x})\big)$ directly from the definition, and the variance derivative is not uniformly signed over $\mathscr{D}$, so the direction of the distortion is calibration-dependent. □

Proof of Theorem 1. Write $\mathscr{W}_k(q) = \mathscr{W}_1(q) - \gamma_k P_k(q)$ with $P_3(q) = K_\xi(q)^2\sigma_\xi^2$ and $P_4(q) = q^2\sigma_\xi^2/V_A$. Then

$$\mathscr{X}_k(q) = \frac{d\mathscr{W}_1}{dq} - \frac{d\mathscr{W}_k}{dq} = \gamma_k\,P_k'(q),$$

so the chain-welfare terms cancel identically and the wedge reduces to the penalty derivative.

For part (i), $P_3'(q) = \frac{d}{dq}[K_\xi(q)^2\sigma_\xi^2] = 2\,K_\xi(q)\,K_\xi'(q)\,\sigma_\xi^2$, with $K_\xi(q) = q\,k^A(q) = q\,\alpha_\theta^A/\Delta(q) > 0$ on $(0,1]$. Differentiating, $K_\xi'(q) = \alpha_\theta^A\,[\Delta(q) - q\Delta'(q)]/\Delta(q)^2$, and a direct expansion of $\Delta(q) - q\Delta'(q)$ from (2) gives the identity

$$\Delta(q) - q\Delta'(q) = \frac{V_A\left(V_C - \delta\,\sigma_\theta^2\right) + \delta^2\,\sigma_\theta^2\,\sigma_\xi^2\,q^2}{V_A\,V_C}. \tag{T1.2}$$

On $\mathscr{D}$ the admissibility bound $\delta < 1 + \sigma_C^2/\sigma_\theta^2$ implies $\delta\,\sigma_\theta^2 < \sigma_\theta^2 + \sigma_C^2 = V_C$, so the numerator of (T1.2) is strictly positive and $K_\xi'(q) > 0$ on $[0,1]$; hence $P_3'(q) > 0$ on $(0,1]$, with magnitude depending on δ through $k^A(q)$ but sign free of it. The identity (T1.2) is verified to machine precision at 1,000 random admissible points in the replication package (`verify_corrected.py`, Test 4).

For part (ii), $P_4'(q) = 2q\sigma_\xi^2/V_A > 0$ for $q \in (0,1]$, and no equilibrium object appears, so the wedge is invariant in δ at fixed q.

For the comparison of optima, both planners share $\mathscr{W}_1$; since $\mathscr{X}_k(q) > 0$ on $(0,1]$, the complete-welfare first-order condition $F^{-1}(q) = d\mathscr{W}_k/dq = d\mathscr{W}_1/dq - \mathscr{X}_k(q)$ crosses strictly below the chain-welfare condition, so $q_k^{soc} < q_1^{soc}$. The decomposition follows from Proposition 3:

$$\Delta\Pi(q) - \frac{d\mathscr{W}_k}{dq} = \left(\frac{d\mathscr{W}_1}{dq} - \Psi(q)\right) - \left(\frac{d\mathscr{W}_1}{dq} - \mathscr{X}_k(q)\right) = \mathscr{X}_k(q) - \Psi(q).$$

Whenever $\Psi(q) \le 0$, in particular under strategic substitutes with $\frac{d}{dq}\mathrm{Var}(\bar{x}) \ge 0$, we have $\Delta\Pi - d\mathscr{W}_k/dq = \mathscr{X}_k - \Psi \ge \mathscr{X}_k > 0$ and over-adoption obtains; the comparison can reverse only when the classical term is positive, which by Proposition 3 requires δ and $\frac{d}{dq}\mathrm{Var}(\bar{x})$ to share a sign, and exceeds the external wedge. □

Proof of Proposition 4. Under the decomposition $\xi = \ell\theta + \zeta$, the platform signal is $\tilde{s}_f^A = (1+\ell)\theta + \zeta + u_f^A$ with variance $V_A(\ell)$. The residual common component of an adopter's forecast error is ζ, so the normalized error covariance of two adopters is $\sigma_\zeta^2/V_A(\ell)$, and the $\mathscr{W}_4$ penalty becomes $q^2\sigma_\zeta^2/V_A(\ell)$. Because $\mathscr{W}_1$ cancels in the wedge,

$$\mathscr{X}_4(q) = \gamma_4\,\frac{d}{dq}\left[q^2\,\frac{\sigma_\zeta^2}{V_A(\ell)}\right] = \frac{2\gamma_4\,q\,\sigma_\zeta^2}{V_A(\ell)},$$

which is strictly positive for $q \in (0,1]$ whenever $\sigma_\zeta^2 > 0$ and vanishes when $\sigma_\zeta^2 = 0$, giving (i).

Holding σ_ξ^2 fixed, ℓ enters only through $V_A(\ell)$, which is increasing in ℓ, giving (ii). No equilibrium loading appears, giving (iii). □

Proof of Proposition 5. Each adopter observes a single processing shock, so individual best responses and hence (k^A, k^C) are as in Proposition 1. Under $\mathscr{W}_3$, the platform-driven component of aggregate variance is $\mathrm{Var}(\sum_j K_{\xi,j}\xi_j) = K_\xi(q)^2\sigma_\xi^2 H$, with H the (share-weighted) concentration index; the penalty derivative scales by H, giving (i). Under $\mathscr{W}_4$, the probability that two random adopters share a platform is $\sum_j (q_j/q)^2 = 1/J$ under symmetry, in which case their errors are perfectly correlated, and with complementary probability they use distinct platforms with correlation ρ; the expected pairwise correlation is therefore $H(J,\rho)$, and the penalty becomes $H\,q^2\sigma_\xi^2/V_A$, giving (ii). Positivity follows from Theorem 1 since $H \geq 1/J > 0$, and $\partial H/\partial J = -(1-\rho)/J^2 < 0$ for $\rho < 1$. □

Proof of Proposition 6. Part (i) is the static variance decomposition applied period by period, using independence of θ_t and ξ_t. For part (ii), θ_t is i.i.d. and independent of ξ at all leads and lags, so

$$\mathrm{Cov}(\bar{x}_t, \bar{x}_{t+1}) = K_\xi(q)^2\,\mathrm{Cov}(\xi_t, \xi_{t+1}) = K_\xi(q)^2\,\phi\,\sigma_\xi^2,$$

by stationarity of the AR(1). For part (iii), their ratio is $\rho_x(q,\phi) = \phi\,R(q)$, with $R(q)$ the variance share defined in the statement. Since $R(q) > 0$ for $q > 0$ and does not depend on ϕ, ρ_x is linear and strictly increasing in ϕ on $(0,1)$.

For monotonicity in q, note that $R(q)$ is strictly increasing in the ratio $K_\xi(q)/K_\theta(q)$. Dividing numerator and denominator by $k^A(q) > 0$ and substituting the closed forms of Proposition 1, the equilibrium denominator $\Delta(q)$ cancels and

$$\frac{K_\xi(q)}{K_\theta(q)} = \frac{q\,\alpha_\theta^A}{g(q)}, \qquad g(q) \equiv q\,\alpha_\theta^A + (1-q)\,\alpha_\theta^C\big(1 - \delta\,q\,\alpha_\xi^A\big).$$

Direct differentiation and simplification give

$$\frac{d}{dq}\left[\frac{K_\xi(q)}{K_\theta(q)}\right] = \frac{\alpha_\theta^A\,\alpha_\theta^C\,\big(V_A - \delta\,q^2\,\sigma_\xi^2\big)}{V_A\,g(q)^2}, \tag{P6.3}$$

which matches the symbolic factorization in the replication package (`verify_dynamic.py`). On $\mathscr{D}$ the bound $\delta < 1 + \sigma_A^2/(\sigma_\theta^2 + \sigma_\xi^2)$ implies $\delta(\sigma_\theta^2 + \sigma_\xi^2) < V_A$, hence $\delta\,q^2\,\sigma_\xi^2 < V_A$ for $q \in [0,1]$,

so the numerator of (P6.3) is strictly positive on $\mathscr{D}$; and $g(q) > 0$ on $[0,1]$ because $\delta\,\alpha_\xi^A < 1$ on $\mathscr{D}$. Hence K_ξ/K_θ, and therefore $R(q)$ and $\rho_x(q,\phi)$, are strictly increasing in q on $(0,1]$. □

Proof of Proposition 7. Part (i): with fundamentals unchanged at τ, the change in aggregate entry is $\Delta\bar{x}_\tau(q) = K_\xi(q)\Delta\xi$ directly from the linear aggregate. Part (ii): the ratio $K_\xi(q_2)/K_\xi(q_1)$ cancels $\Delta\xi$; its monotonicity in q is inherited from $K'_\xi(q) > 0$, established via (T1.2) in the proof of Theorem 1. Since $K_\xi(0) = 0$ and $K'_\xi(q) > 0$ on $\mathscr{D}$, the ratio exceeds one for $q_2 > q_1 > 0$. Part (iii): with (q_1, q_2) and the signal structure known up to scale, the ratio is a known function of δ alone; injectivity of $\delta \mapsto K_\xi(q_2;\delta)/K_\xi(q_1;\delta)$ on $\mathscr{D}$ is verified numerically across the admissible region in the replication package, and point identification follows wherever it holds. Since $K_\xi(0) = 0$ and, by the identity (T1.2) in the proof of Theorem 1, $K'_\xi(q) > 0$ on $\mathscr{D}$, the magnitude $|\Delta\bar{x}_\tau|$ is zero at $q = 0$ and strictly increasing on $(0,1]$. □

Proof of Proposition 9. Conditional on both evaluators of a pair being platform adopters, their forecast errors share the common shock ξ_m, whose contribution to the normalized covariance is σ_ξ^2/V_A; conditional on any other technology profile, the errors share no common component and the covariance is zero. Evaluators are drawn independently from the population, so each of the $\binom{n}{2}$ pairs is a co-adopting pair with probability q^2, giving $S(q) = \binom{n}{2} q^2 \sigma_\xi^2/V_A$ exactly. The wedge is $\gamma_4 S'(q)$ because $\mathscr{W}_1$ cancels, and $S'(q) = \binom{n}{2} 2q\,\sigma_\xi^2/V_A > 0$ for $q \in (0,1]$. No equilibrium loading appears, so the wedge is invariant to δ at fixed q. □

Proof of Proposition 8. Write $q_{t+1} = G(q_t,\lambda) \equiv (1-\lambda)\,q_t + \lambda\,F(\Delta\Pi(q_t))$. The three parts share one engine: the path q_t is increasing in λ (part i), which pushes the per-period wedge $w(q_t)$ up at every date (part ii) and the autocorrelation $\rho_x(q_t,\phi)$ up along the path (part iii).

(i) Since $\delta \geq 0$ implies $\Delta\Pi'(q) \geq 0$ by Proposition 2, and F is strictly increasing, $\partial G/\partial q = (1-\lambda) + \lambda\,F'(\Delta\Pi(q))\,\Delta\Pi'(q) > 0$. Along the transition path $q_t < F(\Delta\Pi(q_t))$, so $\partial G/\partial\lambda = F(\Delta\Pi(q_t)) - q_t > 0$. The chain rule then gives

$$\frac{dq_{t+1}}{d\lambda} = \frac{\partial G}{\partial q}\frac{dq_t}{d\lambda} + \frac{\partial G}{\partial \lambda} > 0$$

under the induction hypothesis $dq_t/d\lambda \geq 0$, with $dq_0/d\lambda = 0$, so q_t is strictly increasing in λ at every $t \geq 1$. Since G is a convex combination of q_t and $F(\Delta\Pi(q_t))$, both in $[q_0, q^*]$, the path remains below q^* and converges to it for every $\lambda \in (0,1]$.

(ii) The per-period wedge $w(q) = 2\gamma_4 q\sigma_\xi^2/V_A$ is strictly increasing in q. By part (i), $q_t(\lambda') > q_t(\lambda)$ for all $t \geq 1$ whenever $\lambda' > \lambda$, so $w(q_t(\lambda')) > w(q_t(\lambda))$ termwise and $\mathscr{L}(\lambda';\beta) > \mathscr{L}(\lambda;\beta)$ for any $\beta \in (0,1)$.

(iii) Along the path, q_t is increasing in t (since $G(q,\lambda) > q$ on $[q_0, q^*)$), and $\rho_x(q,\phi) = \phi R(q)$ is strictly increasing in q by Proposition 6(iii), so $\rho_x(q_t,\phi)$ is strictly increasing in t. □